

Mediated Cognition:
Information Technologies and the Sciences of Mind

Alasdair J. McMillan

A Dissertation submitted to the Faculty of Graduate Studies in Partial Fulfillment of the
Requirements for the Degree of Doctor of Philosophy

Graduate Program in Science and Technology Studies
York University
Toronto, Ontario

August 2016

© Alasdair McMillan, 2016

Abstract

This dissertation investigates the interconnections between minds, media, and the cognitive sciences. It asks what it means for media to have effects upon the mind: do our tools influence the ways that we think? It considers what scientific evidence can be brought to bear on the question: how can we know and measure these effects? Ultimately, it looks to the looping pathways by which science employs technological media in understanding the mind, and the public comes to understand and respond to these scientific discourses. I contend that like human cognition itself, the enterprise of cognitive science is a deeply and distinctively mediated phenomenon. This casts a different light on contemporary debates about whether television, computers, or the Internet are ‘changing our brains,’ for better or for worse. Rather than imagining media effects as befalling a fictive natural mind, I draw on multiple disciplines to situate mind and the sciences thereof as shaped from their origins through interaction with technology. Our task is then to interrogate the forms of cognition and attention fostered by different media, alongside their attendant costs and benefits.

The first chapter positions this dissertation between the fields of media studies and STS, developing a case for the reality of media effects without the implication of ‘technological determinism.’ The second considers the history of technological metaphor in scientific characterizations of the mind. The third section consists of three separate chapters on the history of cognitive science, presenting the core of my case for its uniquely mediated character. Across three distinct eras, what unifies cognitive science is the quest to understand the mind using computational systems, operating by turns as generative metaphors and tangible models. I then evaluate the contemporary cognitive-scientific research on the question of media effects, and the growing role of electronic media in science. My fifth and final section develops a content analysis: what is said in the media about the popular theory that media themselves, in one way or another, are causing attention deficit disorders? The work concludes with a summary and some reflections on mind, culture, technoscience and markets as recursively interwoven causal systems.

Acknowledgements

Though writing is typically a solitary pursuit, I couldn't have done this without a lot of help. Infinite thanks to my wife Danielle for her continuous emotional and material support over many years of writing, and to our daughter Elspeth for being the most wonderful distraction from this work that I could ever hope for. Thanks as well to my parents; your love and support, along with your own senses of intellectual curiosity helped set me on this path and get me to the end of it.

A huge thank you to my supervisor, Edward Jones-Imhotep, for your constructive advice and your imperturbable calm; thanks as well to my committee members Chris Green and Mike Pettit. All three of you struck just the right balance between critique, appreciation, and "get it done!" I'm also very grateful for constructive feedback from and discussions with my external examiners, Steve Bailey and Tara Abraham.

Thanks, too, to the graduate program directors over my time here and since the birth of the program: Bernard Lightman, Kenton Kroker, and Aryn Martin. Each of you helped shepherd me through and make it a fulfilling experience, alongside the rest of the STS faculty and my graduate student colleagues. It's been a pleasure getting to know you and sharpening my thinking in discussions at seminars, lectures, and conferences. Foremost among colleagues and friends in the program I want to thank Cameron Murray and Bernhard Isopp, without you guys I would have had nowhere near as much fun doing this.

Acknowledgements are also due to the nonhuman actors deployed in writing the text, both hardware and software. Without Zotero there is no way my puny human brain could have wrangled all these references, thanks be to it and the volunteer team of human contributors to this open-source project. And I may have had to swap out a few components here and there but never once lost any substantial amount of work, nor had to give a second thought to backups. Call it a regression to the days of mainframes and timesharing if you will, but I say cloud storage is a wonderful thing.

Final acknowledgements are in order to the Social Sciences and Humanities Research Council of Canada, whose grant support was essential in allowing me to complete this research project.

Table of Contents

Abstract	ii
Acknowledgements.....	iii
Table of Contents	iv
Table of Figures	v
Prelude.....	1
1. Messages and mediators.....	7
The hazards of media effects.....	10
Mediation, ANT and STS.	21
The sciences and economies of attention.	34
2. Technological metaphor and the figuration of mind.....	45
3. A mediated history of cognitive science	75
3.1 : Cybernetics: electromechanical brains and analog minds.....	84
A technoscientific metaphysics.	84
From ordnance to Perceptrons: thinking machines in cybernetics.	90
Cybernetics’ speculative futures.	118
3.2 : Cognitivism: the ascent of digital representations	130
The emancipation of models and the return to minds.....	130
Programming tools for cognition: Logic Theorist, General Problem Solver, and list processing.	138
Cognitivism becomes ‘old-fashioned.’.....	166
3.3: Post-Cognitivism: PDP, Connectionism, and a return to neurobiology	178
4. Media effects in science and society.....	216
Delimiting ‘effects.’	220
Toward neuromediation.....	231
Another angle on mediation in science and cognition.	265
5. Technology and attentional pathology: the case of ADHD	288
5.1. Historicizing diagnosis.....	292
Regarding ‘construction.’.....	292
From ‘moral treatment’ to ‘moral control.’.....	298
Hyperkinesis, psychoanalysis, and psychopharmacology.....	303
Regimes of treatment, regimes of diagnosis.	310
5.2 : Media research.....	318
Method.	318
Quantitative findings.....	325
Qualitative analysis.....	339
Epilogue	377
Bibliography.....	401
Appendix: Figure references and permissions	427

Table of Figures

Figure 1 - Visual analogy between spine and telephone system	67
Figure 2 – Constituent disciplines of cognitive science	77
Figure 3 - ‘Cybernetic Séance,’ manipulated composition by Suzanne Treister	92
Figures 4-8 - W. Ross Ashby's Homeostat	103
Figure 9 – DAMS.....	106
Figure 10 - Grey Walter’s tortoise and ‘cyborg family’	108
Figure 11 – Perceptron block diagram.....	111
Figure 12 - Organization of the Mark 1 Perceptron, from operator manual.....	112
Figure 13 - Cover of Minsky and Papert’s Perceptrons	116
Figure 14 - ARPANET Logical map.....	134
Figure 15 - Logic Theorist sketch.....	149
Figure 16 - List programming	152
Figure 17 - Planning diagram.....	157
Figure 18 - SHRDLU’s World	160
Figure 19 - HARPY Block diagram.....	167
Figure 20 - Copycat table of results.....	184
Figure 21 – Rodney Brooks’ subsumption architecture	187
Figure 22 – Henry Markram’s path of progress.....	198
Figure 23 - Human Brain project graphic.....	199
Figure 24 - Promotional graphic for ‘neuromorphic’ computing, IBM Research.....	207
Figure 25 - SPAUN and the ‘Serial WM’ test.....	209
Figure 26 - Cultivation analysis summary	225
Figure 27 - Brain images of media violence exposure.....	238
Figure 28 – Dimitri Christakis’ ‘mouse overstimulation chamber’.....	247
Figure 29 - Lumosity performance tracking screenshot	254
Figure 30 - Eyewire screenshot.....	274
Figure 31 - DARPA SUBNETS technology rendering	281
Figure 32 - Bensedrine study results	307
Figure 33 - Media sampling queries.....	319
Figure 34 - Coding exemplars for neuroessentialism	323
Figure 35 - Prominent keywords in corpus (with total occurrences).....	330
Figure 36 - Breakdown of sample by year	332
Figure 37 - Lorenz attractor.....	345

Prelude

This project originated with my interest in an increasingly popular opinion: that new media are changing the way that we think. I routinely hear this view expressed in conversations with friends, family, and colleagues. It is also the topic of much discussion *within* the media. There are of course many narratives about how new technologies, particularly computers and the Internet, might make us more intelligent. In their most current formulations, our collective intelligence is poised to be exponentially augmented through interaction with distributed, networked machine learning systems (whether or not we wish to call them ‘artificial intelligences;’ Kelly, 2014). Equally, however, we find pieces tending toward the opposite view, often headlined with troubling questions. Is Google making us stupid, for instance, rather than augmenting our cognition? In a thusly-titled essay, Nicholas Carr compares himself to HAL 9000 at the end of *2001: A Space Odyssey*, having the circuits which make up his mind disconnected one by one: “Over the past few years I’ve had an uncomfortable sense that someone, or something, has been tinkering with my brain, remapping the neural circuitry, reprogramming the memory. My mind isn’t going—so far as I can tell—but it’s changing. I’m not thinking the way I used to think” (Carr, 2008). Google is merely a placeholder, as he really blames the whole information-foraging, hyperlink-oriented experience of the Internet for his inability to engage in deep, sustained attention and thought. For Carr, this medium is making us not only stupid but shallow people (Carr, 2011).

Another recent piece asks if the social network Facebook could be ‘making us lonely,’ invoking the case of Yvette Vickers, “a former *Playboy* playmate and B-movie star, best known for her role in *Attack of the 50 Foot Woman*,” who died alone and was not found for nearly a year, her

body “mummified, near a heater that was still running. Her computer was on too, its glow permeating the empty space” (Marche, 2012). Without children or an “immediate social circle of any kind,” she had begun to look elsewhere for companionship, and leading up to her death had made calls “not to friends or family but to distant fans who had found her through fan conventions and Internet sites” (*ibid.*). The implication is clearly that none of those online connections were robust enough for anyone to check on her in the real world, despite her death becoming within weeks the subject of Facebook posts numbering in the thousands. The author refers to survey research suggesting that heavy users of the site, despite communicating with perhaps hundreds or thousands of ‘friends’ online, nevertheless paradoxically scored higher on measures of loneliness and depression. As one amusingly titled study asks, could it be that Facebook is “creating iDisorders?” (Rosen, Whaling, Rab, Carrier, & Cheever, 2013). This and other studies do lend limited support to the idea, though really what they describe is a correlation of anxious and depressive symptoms with excessive usage of social networking sites, alongside texting and other forms of mediated communication (see also Clayton, Osborne, Miller, & Oberle, 2013; Steers, Wickham, & Acitelli, 2014). If the Internet is changing how we think, it would be doing so largely by changing how we interact.

Scientific research is often recruited in such debates, as is talk of the brain. There is a good deal of survey research now examining psychological correlates of Internet use. Despite all the brain-talk, however, actual neuroscientific research into media effects is rather sparse. As Steven Pinker notes, “critics of new media sometimes use science itself to press their case, citing research that shows how ‘experience can change the brain.’ But cognitive neuroscientists roll their eyes at such talk. Yes, every time we learn a fact or skill the wiring of the brain changes; it’s not as if the information is stored in the pancreas” (Pinker, 2010). He stands firmly in the cognitive augmentation

camp, however, contending that our improved capacities for managing, searching and manipulating our ‘collective intellectual output’ can only be a good thing, and that “far from making us stupid, these technologies are the only things that will keep us smart” (*ibid.*). In the most general sense, this dissertation is about these debates. It is a multifaceted exploration of this idea that new technologies may be transforming human thought. I am concerned throughout with media effects as Marshall McLuhan framed them, “the psychic and social consequences of [their] designs or patterns as they amplify or accelerate existing processes” (McLuhan, 1994, p. 8). In what ways might media reshape consciousness, perception, and cognition? The question remains vital, as technology seems to be changing and proliferating more rapidly than ever. The media landscape has been wholly reconfigured since the time of McLuhan’s *Understanding Media*, and yet some of his proclamations seem to make more sense in the age of the Internet than they must have in 1964 (Lapham in McLuhan, 1994, p.x). Inspired by mediated narratives reflecting upon the effects of new media, this work sets out to examine the complex histories and contemporary realities of reciprocal interaction between science, media, and minds.

My goal is to understand the genesis of media effects discourse through a looping series of exchanges between these domains. Beyond the effects themselves, I consider how we come to know them. This inquiry is driven by three foundational, interrelated questions: how do the material and theoretical practices of cognitive scientists enact and construct ‘media effects;’ how does their research affect popular conceptions of how the mind works and what cognition ‘really is;’ and what, in turn, are the broader effects of our conceptions of the cognitive? I wish to simultaneously steer away from and investigate the common popular framing of the issue in terms of harm in child development (Anderson & Bushman, 2001; Carnagey, Anderson, & Bushman, 2007; P. Greenfield & Yan, 2006), instead adopting the hypothesis that cognition has always depended upon and

evolved through embodied interactions with external media (Clark, 2003; Hutchins, 1996). This position is closely aligned with McLuhan's understanding of media, going further still in emphasizing that media are not mere 'extensions of man,' but internal to human nature itself. If we are 'hard-wired' for anything, it is to be malleable organisms, capable of extending our minds and capacities through tool-use. I seek to engage more fully with scientific research into these questions than McLuhan and many others within the media studies tradition have done. My purpose, ultimately, is not to conflate legitimate and important research in this field with alarmist or otherwise distorted reports in the popular media. Rather, I intend to carefully consider the relationship (or occasional lack thereof) between mediated popularization and specialist scientific discourse.

This project is focused closely on cognitive science. I situate this interdisciplinary field as distinct from but intertwined with psychology, neuroscience, and computer science. It is of particular relevance to my project for its own mediated nature, which I expound in three historical chapters at the heart of this text. Cognitive science seems poised to clarify the question of media effects, perhaps to move beyond evocative metaphors and to access the hard facts of the matter. At the same time, however, it is animated throughout its history by a particularly tenacious technological metaphor: that of mind as computation. It seeks to understand the mind by using computer technology as metaphor and model. Why is this tool so central to the field, and what implications does it have for the character of its research? Scientific research, I find, is best seen not as a fount of neutral and objective information on the cognitive effects of media, but as itself expanding the frontiers of mediated cognition. Laboratories are heterogeneous spaces where minds link together in networks with nonhuman organisms and apparatuses, both technological and institutional. Accounts which fail to critically assess the social and cultural dimensions of scientific knowledge production remain fundamentally incomplete, doing as little justice to the question of

media effects as those which ignore cognitive science altogether. Through this analysis I wish to contribute to understanding the developing culture of the cognitive brain, with culture understood as “the images which make imagination possible, in the media with which we mediate experience” (Strathern, 1992, p. 33). How do scientific conceptions of cognition circulate in mediated images, and how then do we collectively imagine the relationship between mind and technology?

Outlining briefly what is to come, the first chapter further elaborates the premises of my project and reviews relevant literature on the concept of media. I contend that the questions of media studies and science studies are as inseparable from one another as are science and media themselves. I also situate attention as another key theme running through the project. The second chapter focuses on technological metaphor, offering my account of its role within science and a history of analogies between technologies and mind. The telegraph is of particular importance in this regard. The following three chapters constitute the aforementioned study of cognitive science, presenting its history as profoundly mediated, and rooted in the metaphor of mind as computation. The three separate chapters each focus on one of the eras I distinguish: first that of cybernetics, followed by that of ‘orthodox’ symbolic cognitivism and artificial intelligence research, then lastly the recent period of renewed interest in connectionism and neurobiology. In each I highlight a few human actors, but just as importantly some specific, paradigmatic technologies used as models for the mind. The sixth chapter is the one most closely focused on the science of media effects, examining the research which has been done within cognitive science and its associated disciplines into the consequences of our interactions with technology. I find a wealth of psychological research into violence, but sparser and more ambiguous evidence on other topics, and few neuroscientific studies addressing these questions. In other ways, however, I find technological mediation is everywhere within the contemporary sciences of mind. The final two chapters then constitute a case

study and media analysis concerning one specific mental pathology, attention deficit disorder. This is often seen as linked with the increasing ubiquity of technology, and so after considering the history of the diagnosis, I undertake a quantitative and qualitative content analysis of published articles discussing this causal claim.

My method is by turns historical, sociological, and philosophical, and I draw largely on the synthesis of materials previously published in one way or another. The same new media I investigate offer a wealth of possible sources, including online journals, correspondence, and discussion fora, rendering public a type of discourse which was once confined to informal and face-to-face interaction, accessible only to direct ethnographic observation. Along with digitized archival and secondary sources, these sites, and the transformations of the public sphere which they have wrought, constitute the virtual 'field' in which I have immersed myself for several years. Relatively few volumes of the traditional paper codex type were handled in the course of producing this document. My digital library in the 'cloud,' and the computational actor to whom I delegated the task of managing it, were of far greater importance. I cannot help but enact the cognitive metamorphoses which I set out to investigate. This is a work of and about the mind in the era of Google, a series of interconnected studies rather than a wholly linear argument. I offer a multiplicity of perspectives on the fundamental questions at stake, developing thorough descriptions without suppressing heterogeneity or uncertainty. I pursue no totalizing, reductive explanations or immutable conclusions, and I harbour no fears about the corruption of our 'natural' ways of thinking. Instead, I call for openness to the multiplicity of causal influences and artificial mediations through which human subjectivity is, and always has been, produced.

1

Messages and mediators

“In the long run, for such media or macromyths as the phonetic alphabet, printing, photography, the movie, the telegraph, the telephone, radio, and television, the social action of these forms is also, in the fullest sense, their message or meaning.”

— Marshall McLuhan (1959, p. 340)

This is an inquiry into minds and media. My aim is to describe and understand some of the many interconnections between shifts in communications technologies and shifts in how we think. How do we employ technologies to extend and restructure our spheres of attention and influence? In what new ways do we see and act through screens, and in what ways do they act upon us? How are these effects of media operative within science, and how do mass media then translate the scientific discourse on media effects? These are among the overarching questions for which I seek answers in this work, and later in this chapter I situate cognitive science as simultaneously a resource and area of inquiry. First, however, I must consider some questions of conceptual vocabulary and methodology, even if in some sense I will be positioning myself against definition and against method.

In the linear medium of print, it is imperative to define one's terms at the outset, and yet the work of definition often yields controversies and tensions which stand unresolved even at the end of the work. The meaning of 'mind' is of course one of the longest-standing and most active topics in philosophy. 'Media' and 'mediation' are concepts with equally multifarious meanings. Both in common usage and in a range of specialized senses, these notions have been articulated in many

distinct but overlapping ways. Alongside their popular usage, this chapter explores these concepts in the senses proposed by Marshall McLuhan, actor-network theory, and information theory. Tracing out these definitions, one observes a vast expansion of meaning. We pass from a narrow subset of communications devices popularly called ‘the media,’ to mediation as a fundamental, quasi-universal principle, binding together the stuff of our mental and cultural worlds.¹ Each sense of the concept is operative within this work.

What follows is at once a project of ‘communications studies,’ of ‘technology studies,’ and of ‘science studies;’ it is the concept of media, however, which reveals an underlying unity of purpose across these disciplines. We tend to apply the term to specific material artifacts employed in communication, and particularly the newest of those. We may include books within our definition, and certainly newspapers, but the concept often calls to mind more recent developments – telephony, television, or the Internet. While I try to sustain a longer view, I follow this popular usage in focusing much of my attention on newer media. Their effects are both primary matters of public concern, and relatively unexamined within science and technology studies. I argue that a broader understanding of media and mediation places these concepts at the intersection of communications studies and STS, crucial to resolving long-standing questions within each discipline. While literature and to some extent journalism already had their own established places in the academy and in society, it was the rise of new electronic media that led Marshall McLuhan to depart from his literary training in developing his own broader concept of media. I contend in this chapter that by understanding any ‘extension’ of human capacities as a medium – not only books and television, but everything from electric light to money to automation – McLuhan’s account of media is closely

¹ This broader definition is linked simultaneously to another, earlier sense of medium, in the way that water and air constitute media, and to the sense proposed by information theory, of the medium as any substrate for a code (thus including not only electromagnetic radiation, but also things like DNA and the fields of fundamental physics).

comparable to those of technology or ‘technique’ within STS and history of science. Across the text which follows, I go on to describe some of the central roles played by electronic media in the constitution of new scientific theories, disciplines, and modes of authoritative demonstration.

Media are the technologies through which communication and information are effected, affected, and understood. They encode and transport sensations and ideas, not only reshaping our possibilities for thought and action, but serving as conceptual paradigms for the very ideas of information, communication, and cognition.² Technologies like the telegraph or the computer may serve as ‘popularizing’ metaphors for the mind, relevant and worthy of consideration in their own right, but they are also often conceptual devices within the practice of science. As Laura Otis suggests, such metaphors may not simply ‘express’ the ideas of scientists, but constitute them. Metaphors “suggest new visions, images and models; they inspire scientists to approach problems in new ways” (Otis, 2001, p. 59). As such metaphorical redescriptions of the mind become solidified within scientific practice, they may gradually cease to be seen as metaphorical and instead as reflecting a more fundamental homology of structure.

Science stands in an ongoing reciprocal relationship with media. Though there is good reason to reject the traditional deterministic narrative, encapsulated by the famous motto of the 1933 World’s Fair – ‘Science Finds, Industry Applies, Man Conforms’ – scientific research has nevertheless long fed in to the development of new communications technologies and their accompanying industries. It is less that an already-constituted entity called ‘science’ produced discoveries which were then applied to the creation of technologies, and more that science,

² There is hence a hidden depth in the phrase ‘information and communication technologies,’ that common institutional shorthand for computational media. The phrase applies just as readily to anything we might place under the general category of medium—old or new, analog or digital. Yet I follow Leo Marx in understanding ‘technology’ as always involving an entire sociotechnical assemblage, rather than simply material artifacts, as in the ‘hazardous’ common view (Marx, 1997). I explore this idea further in relation to debates around ‘technological determinism’ below.

technology, and communication all started to take their modern shape in the context of the most significant media innovation since the alphabet: the electric telegraph, the first means by which “messages could travel faster than a messenger” (McLuhan, 1994, p. 89). While other factors were doubtless operative in this transition, the rise of a global telegraphic network was deeply entwined with the shift from an era of natural philosophy and the mechanical arts to the institutional, industrial modes of science and technology we know today.³ As I discuss at greater length in the next chapter, the telegraph also served as a crucial technological metaphor within the nascent neurosciences of the era (Lenoir, 1986, 1994; Otis, 2001). While telegraphy in itself did not cause a scientific revolution, it stands as an epoch-making symbol for the co-constitution of science, technology, and communications media.

The hazards of media effects.

In contemporary Western culture and its global outposts, the media whose effects are primary ‘matters of concern’ (Latour, 2004b) are television, computer systems, and the Internet. Perhaps more than any of them individually, we also worry about the increasing ubiquity of all these media forms. The aim of this project is to follow the trajectory of this matter of concern as it passes through its own mediated circuits, from scientific laboratories through to popular discourse. Following the terms of these debates, my analysis will often focus on these novel technologies. I also wish to emphasize, however, the sense in which these concerns are not so new, and a broader view

³ The terms ‘scientist’ and ‘technology’ were both coined in the nineteenth century, by William Whewell in 1833 and Jacob Bigelow in 1828, respectively, but did not come into common usage until the twentieth century (Marx, 1997; Yeo, 2003). Other factors operative in this transition include the related development of other sociotechnical systems such as the railroad, and a range of interconnected cultural, political, and economic changes.

of mediation. The kinds of hand-wringing press narratives which I discussed in the Introduction are only the most recent manifestation of an age-old fear about the effects of communications media on the mind and on society. Plato's mythic narrative of the origins of writing in the *Phaedrus*, and Jacques Derrida's deconstruction thereof, are sufficiently well-known as to scarcely bear repeating. Nevertheless, they set the stage for all that follows. The Egyptian god Theuth, mythic inventor of 'many arts' but above all written letters, boasts that "this discipline (*to mathēma*), my King, will make the Egyptians wiser and will improve their memories: my invention is a recipe (*pharmakon*) for both memory and wisdom" (in Derrida, 1981, p. 75). Derrida's deconstruction emphasizes the double valence of this term *pharmakon*, as both remedy and poison: precisely elided by one common translation of the term as 'a specific,' but made clear in the response of Plato's Thamus.

"This discovery of yours will create forgetfulness in the learners' souls, because they will not use their memories; they will trust to the external written characters and not remember of themselves" (Plato, *Phaedrus*): with this warning, the mythic god-king heralds the advent of recursive media critique, that writing of fears about writing which has now metamorphosed into new forms.⁴ Now this section of *Phaedrus* is cited in discussions of whether the Internet could be 'rewiring our brains' (Mills, 2014). As with many other aspects of Platonism, however, the privileging of natural speech must be overturned. While writing may have retained some novelty in his time, by now we can recognize there is nothing more essential to human cognition than entrusting portions of our knowledge to external media. Without such mediation, what we recognize as 'rational thought' and science would be impossible. Hence I adopt McLuhan's broader concept of medium, referring to

⁴ This is, moreover, far from an isolated instance of this concern within the Platonic *oeuvre*. His imperative to control or exile the storytellers in the *Republic* may likewise be read as another affirmation of media effects, and he places a surprising emphasis on the proper forms of musical expression conducive to a well-ordered state. The modes of Greek music are systematically evaluated in the *Republic* in terms of whether they are acceptable or not for the ideal state, and which groups should be permitted to listen to each, be they men or women, soldiers or legislators.

any extension of human perceptual or behavioral capacities. In this sense, spoken language and writing are as much media as television or the Internet, as are even less obviously ‘communicative’ technologies like roads, electric lights, and money. The significance of McLuhan’s best-known aphorism, ‘the medium is the message,’ is not that the explicit content of communication is wholly insignificant, but that it operates at a higher conceptual and cognitive level which conceals the reshaping of perception and sensation by the *forms* of communication. For the most part, existing studies in the public understanding of science, or science and the media (e.g. Nelkin, 1995; Bucchi, 1998), focus their attention on how the content of scientific research is communicated through mass media. While I do not seek to dismiss or bracket content as definitively as McLuhan, I wish to bring his emphasis on form into dialogue with this tradition. *What* is communicated in the media, on this view, is a concern secondary to the cognitive and social consequences of *how* communication is carried out.

It may be jarring to see writing, or numbers, or money referred to as media. Yet they are as fundamentally communication-oriented as television and other more modern technologies. McLuhan would argue that it is precisely because old media have become familiar, routinized extensions of our sensorium that they are no longer so readily perceived as ‘media.’⁵ The most effective and potent media, on his account, are those like the electric light – pure medium without content, whose real message is found in the countless domains of human conduct reshaped by this

⁵ Media theorist Vilém Flusser expresses a similar thought in different terms: “People are not always fully conscious of the artificial character of human communication—the fact that man makes himself understood through artistic techniques. After learning a code, we have a tendency to forget its artificiality. If one has learned the code of gestures, then one no longer recognizes that head nodding signifies ‘yes’ only to those who make use of this code” (Flusser, 2002, p. 3). Equally, to those well versed in the codes of scientific representation and simulation, they may cease to be perceived as mediating technologies and instead as direct conduits to reality or Nature. Part of this project’s goal is to explore the implications of a simple but significant idea in one key domain of scientific practice: that the technologies used to produce theories are media, not fundamentally different in kind from those used to communicate and ‘popularize’ them.

mode of artificial illumination. Or written language: a medium so profoundly internalized that its linear mode is inextricable from the constitution of literate, rational subjectivity. It is not merely the familiarity of old media which makes them seem less worrisome, but rather that their characteristics have shaped the baseline normality against which change is measured. Plato's Thamus laments the decay of an oral culture structured around internalized memory (that which Bernard Stiegler calls 'primary retention'), just as more recently public fears centred on the displacement of literate rationality by the mass culture of radio and television, a contest for attention between competing modes of external 'secondary retention' (Stiegler, 2010). I am particularly interested in new media not merely for their novelty, but for the kinds of hybrid forms nurtured by digital technology. McLuhan insisted that the content of a medium was always another medium, and this is surely true of the Internet, which seems less a medium with its own proper form than a *milieu* in which all other extant media can be encoded and interlinked. Yet investigating these effects demands that we sustain a long view, as concerns about the effects of new media may reveal the shaping of cognition and society by older forms, while initial responses to these older forms were equally revealing.

Technological metaphor is the most important concept guiding the following analysis. I mean this first in the common sense routinely employed by practitioners of STS (Edge, 1974; Otis, 2002). Across the next chapters, I examine the different ways in which technological forms like the telegraph or the Turing/von Neumann architecture of computing have been deployed as conceptual redescriptions of the mind. However, McLuhan also points toward a deeper sense of this idea, by way of the term's etymology:

The word 'metaphor' is from the Greek *meta* plus *pherein*, to carry across or transport. In this book we are concerned with all forms of transport of goods and information, both as metaphor and exchange. Each form of transport not only carries, but translates and transforms, the sender, the receiver, and the message. The use of any kind of medium or extension of man alters the patterns of interdependence among people, as it alters the ratios among our senses (McLuhan, 1994, pp. 89–90).

This description of his project highlights not only a more profound conception of metaphor, but also the principal tension within McLuhan's concept of media. A medium may be an extension of human sensation, but it is not a *mere* extension, nor just a matter of sensory experience. The primary effects of media occur at a sensory level, in terms of the 'ratios' of between the senses which they produce. His categories of 'hot' and 'cool' media are directed toward these ratios: cool media like speech or television involve multiple senses, lower definition, and greater participation by the recipient in filling in the details; hot media, like print or the radio, transmit a higher definition of information through one sensory channel. As media of differing character take hold in a society, both individual subjects and the character of their interactions are reshaped. Media are thus metaphors in many senses. They do not just extend but restructure our sensory capacities, and their effects are equally cognitive, social, and political.⁶

No less than any other domain of human interaction, science is profoundly affected by communications media. Harold Innis described these influences in terms of temporal or spatial biases, inhering in the forms taken by communication and their physical substrates. Temporally-biased media like oral speech, stone tablets, or parchment are durable, but difficult to reproduce or transport across long distances; they privilege continuity with the past, and mythic, religious traditions. Spatially-biased media like papyrus, paper, and print are fragile, but conducive to long-distance transportation, thus allied with secular authority, empire, and philosophical or scientific

⁶ James Carey, not generally sympathetic to McLuhan's approach, contrasts McLuhan's work with that of his teacher, Harold Innis, along these lines, arguing that Innis is focused more on the social impacts of communications technologies, and McLuhan more on the cognitive. He acknowledges that both do occasionally discuss each dimension, but that McLuhan unduly privileges the sensory and mental, systematically conflating it with the social (Carey, 1967). While I agree with his characterization of the two authors' respective emphases, many passages like the one cited above indicate that it is mistaken to argue that McLuhan pays insufficient attention to the social. I return to this argument later in this section, where I make the case that the idea of distributed cognition undermines Carey's view of the cognitive and social as truly separate domains.

rationality.⁷ Cultural, social, and economic history for Innis is seen in terms of a continual oscillation between these two biases, with the ideal political entities those best able to counterbalance one with the other: the ancient Egyptians, the network of Greek city-states, and the Roman Empire. The origins of science, on his view, lay in the cross-fertilization of a decentralized Hellenic oral tradition with nascent imported techniques for recording dialectical and spatial demonstration – the alphabet and geometry – while its diffusion and development across Europe were driven by cheap supplies of papyrus and reliable routes of transportation within the Roman Empire (Innis, 2007 [1950]).

McLuhan focused less on the transitions and social negotiations between orality and literacy, instead centering his analysis on the shifts from manuscript culture to print, and the ongoing shift to ‘electric’ media. He offered suggestive indications that the “principles of continuity, uniformity, and repeatability” inherent to the printed word were in turn the basis of modern scientific thought (and equally of nationalism and industrial production: McLuhan, 1994, p. 77). The disposition of scientific reason toward ‘efficient causality,’ mechanistic sequences of cause and effect, and uniform natural laws of infinite extension is, he implies, borne of the linear, replicable organization of print. Likewise, print gives rise to new communications channels and discursive networks of scholars communicating with one another, newly empowered to cite, comment upon, and critique a shared corpus of literature. McLuhan’s speculations inspired a more rigorous historical treatment by Elizabeth Eisenstein, who took up scientific change as a key component within the spectrum of

⁷ For most of recorded history, of course, there was no recognizable separation between philosophical and scientific rationality. I cannot fully explore the implications of Innis’ communications theory here, and the categories of temporal and spatial bias are not central to my analysis, since the hybrid media forms typical to the digital era belie easy categorization into either domain. In a sense, Innis’ contention that the ecology of contemporary media skews toward spatial bias holds true, yet it seems that spatially biased systems like those of computerized global logistics and the Internet may equally serve as ‘support systems’ for the transportation and translation of temporally biased forms, especially in realms of scholarly inquiry: as in the ‘cold chain’ that ensures the safe transportation of genetic materials between laboratories, or the digital archiving of oral history. Deeper analysis of scientific mediation through Innis’ categories certainly merits further investigation, but stands outside the scope of this project.

social transformations wrought by the printing press. The ferment of the Copernican revolution, she contends, had little to do with any remarkable changes in the data available to scientists (*cf.* Kuhn, 1957), and far more with the consequences of print for preserving, duplicating, and disseminating that data.

Conventional narratives of the turn from ‘the little books of men’ to the ‘Great Book of Nature’ miss this crucial point. Access to the Book of Nature did not fundamentally change in the course of the Copernican Revolution, but both access to and the character of the ‘little books of men’ did. It was not so much that print gave rise to the drive for precision and uniformity in knowledge of the world, which had certainly always existed; rather, as Eisenstein puts it, “before the advent of printing, to call for more precise and uniform standards was to indulge in wishful thinking rather than contribute to research” (Eisenstein, 1979, p. 468). Accurate reproduction of textual information was important, but she accorded still greater significance to the faithful replication of maps, charts, and tables, whose engraved forms alongside the text were far more useful than imprecise hand-drawn manuscripts. Eisenstein is put off by McLuhan’s bombastic and often poorly-documented style, but she constructs a serious historical analysis which is fundamentally in line with his conclusions: that the revolutions in science and culture across the early modern period cannot be explained by any radically new data about the world, nor by any transcendent individual genius, and are tied, rather, to the quiet collective transformations wrought by printed media.

McLuhan suggests in his later work that the twentieth century witnessed a parallel shift in the transition to ‘electric’ media: within science and the wider culture, the bias for linear organization and ‘efficient causality’ was subverted, as absolute Newtonian space and mechanistic cause-and-effect were progressively supplanted by relativistic and probabilistic understandings. He draws a remarkable, evocative parallel, one which finds support not only in debates over relativity and

quantum mechanics, but also in relation to the broader turn toward chaos and nonlinearity that exploded into popular consciousness simultaneously with the rise of the Internet (Prigogine & Stengers, 1984).⁸ This sort of claim also indicates the difficulties in working with McLuhan's ideas, however. His claims can't be evaluated, he argues, by the conventional standards ingrained through print: to look for some kind of archival proof or causal sequence whereby we might contend that Einstein or Heisenberg were consciously inspired by an epochal shift to electric communication would be to implicitly support an anachronistic print-derived model of linear causality.

Logical sequence and explanatory coherence, on this view, are not trans-historical universals of rationality, but products of print culture which can only now be recognized as such in light of new electric forms. In his *Gutenberg Galaxy*, he draws conceptual inspiration from Innis, with 'galaxy' connoting an alternative to sequential argument, the setting-up of a 'mosaic' or 'field' which might better serve as guide to understanding a set of complex, multidirectional interactions. The very form of the book is oriented toward linearity, of course, but McLuhan viewed himself as following Innis in attempting to analyse changes in media by pushing the envelope of an old medium. He insisted his claims were best evaluated by their efficacy as 'probes,' radical and performative statements intended to provoke new kinds of awareness about the media, and draw out hidden effects of changes in the dominant forms thereof.

Such probes are aphoristic rather than propositional, akin to the oral tradition which print replaced, and toward which he believed electric media were shifting us back; as explorations rather than explanations, they constitute in his terms a cooling-down of the hot medium of print, pressing

⁸ Likewise, I contend this view on causality is analogous to the same philosophical movements which Prigogine and Stengers sought to ally with nonlinear thermodynamics, namely the French school of philosophy deriving from the likes of Deleuze and Derrida, which sought to privilege the play of difference over the formation of stable identities, as well as Latour and Michel Callon's rethinking of the sociology in terms of 'description' rather than 'explanation.' These parallels may be seen as a kernel of truth within assorted poorly-conceived arguments that McLuhan was a forerunner of so-called 'postmodernism.'

toward greater participation in depth from within a medium which implies a high volume of information in a single sensory mode.⁹ McLuhan's probes inspire the overall parameters of my inquiry, offering a set of conceptual tools for describing and understanding, not an ensemble of causal arguments to be evaluated and applied. Partly as a matter of my own taste, and partly as a reflection of the institutional imperatives of the dissertation form, this project is executed in a 'hotter,' more linear mode than McLuhan favoured, but my conclusions are likewise better described in terms of systems, fields, and mosaics than straightforward causes and effects.¹⁰ Equally, this is in keeping with Latour's exhortation to describe rather than to explain (Latour, 2007). My examination of media effects does not aim to reveal a hidden causal determinant of scientific knowledge, but instead attends to the array of technologies which sustain a discourse of scientific truth about minds and media alike.

This calls to mind the other challenge in extending McLuhan's approach, namely the spectre of 'technological determinism.' His emphasis on the effects of media is often taken to imply a unidirectional shaping of minds by media, and some of his more hyperbolic statements may well be enlisted to support this reading. One way he rephrases the notion that 'the medium is the message,' for instance, is that "the effects of technology do not occur at the level of opinions or concepts, but alter sense ratios or patterns of perception steadily and without any resistance" (McLuhan, 1994, p.

⁹ Along these lines, we may also interpret McLuhan's participation within the mediated culture of his day as equally part of his *oeuvre* alongside his traditional books and more unorthodox publications - *The Medium is the Massage*, his *Den-Line* newsletter (its name deriving from the same nuclear early-warning systems discussed in (Edwards, 1997). Whether appearing on talk shows to offer cryptic *bons mots* ("Marshall McLuhan is taken far too seriously:" (*Marshall McLuhan Interview 1967, 1967*), or to skewer academic reception of his work in *Annie Hall*, he was as much actor as observer with respect to the media of the 1960s and 1970s. Like Bernard Stiegler, whose arguments I address below, McLuhan no doubt regarded some aspects of print-based rationality as worth preserving, and his traditional books do represent attempts to comprehend new media within the linear logic of older forms, yet he's always pushing against the conventions of form, attempting to explore new kinds of rationality more adequate to the age of 'electric' media.

¹⁰ Though I cite them only in passing, the German theorists Friedrich Kittler and Vilém Flusser constitute further background influences for my approach to media, and I situate my contributions here in dialogue with these 'philosophies of media,' along with the work of Derrida and Stiegler.

18). Likewise he regularly describes the content of media as a mere distraction or epiphenomenon, the ‘juicy meat’ which distracts the ‘watchdog of the mind’ and allows the passive, unconscious sensorium to be reshaped without conscious awareness or opposition (*ibid.*, p.32). Taken as straightforward factual claims, these are dubious in the extreme, and indeed suggest an unsustainable technological determinism. I read these instead as hyperbolic probes, aimed at countering a common and equally unsustainable social determinism or instrumentalism, a view which holds that the consciously available content of a given medium is its only relevant feature, and that the effects of technologies are exclusively a matter of the ends to which humans choose to apply them.

Those who wish to find evidence of determinism within McLuhan’s writings will have no trouble. But conversely, the very aim of his work in commenting upon new media is to suggest that society *can* shape technologies in accordance with collective goals, provided that their modes of action are properly understood: “there is absolutely no inevitability, so long as there is a willingness to contemplate what is happening” (McLuhan & Fiore, 2001 [1967]). Along similar lines, McLuhan insisted that the effects of a technology were not solely determined by its intrinsic characteristics, but developed in a dialectical relationship with the existing character of a society.¹¹ Thus the advent of television and ‘electric’ interconnectedness constituted a sharp transition back toward cool oral/aural culture for the print-oriented West, and a radical shock to the system, while the effects of radio in cultures that had never been so deeply shaped by print would necessarily be quite different (McLuhan, 1994). Media effects, once properly understood, could be harnessed and redirected. The

¹¹ The reference to dialectic here suggests a parallel with (Karl) Marx that is not accidental. McLuhan’s stance toward the ‘determination’ of social interaction through technology is similar in character (though not in the implications drawn) to that of Marx, who is also routinely mistaken for a simple determinist. As Donald Mackenzie has persuasively argued, Marx viewed the ‘superstructure’ of social order as determined by the ‘factors of production,’ but defined the latter term far more broadly than technology, to include labourers and a whole associated sociotechnical assemblage (MacKenzie, 1998). Both view humans and machines as interacting in complex ways, with machines sometimes shaping human action, sometimes shaped by humans, and sometimes shaped toward bringing about specific patterns of social organization.

specific conclusions he derived from that were somewhat dubious – suggesting, for instance, that cultured might be ‘programmed’ and directed in a sense by shifting the balance of hot and cool media – but are less significant for my purposes than his overall view of the relationship between technology and society, which is more of the reciprocal and feedback type than any unidirectional determinism.

McLuhan’s technological determinism is thus of a rather ‘soft’ variety. Technology is shaped by humans to extend our capacities in some ways, rather than developing according to an autonomous internal logic, and though media have effects for both minds and societies, these are not *determinate*, linear and universal effects, but contingent and contextual effects. They manifest by an interactive process within a given cultural context, itself shaped by a past history of mediation and its own local variations. A useful parallel within the history of technology literature is Leo Marx’s ‘Technology: The Emergence of a Hazardous Concept’ (Marx, 1997). Noting the seeming ubiquity of technology in popular history and anthropology as a driver of change, and of notions that ‘technology is changing our lives,’ he reminds us that the term itself is a relatively novel coinage: of the same 19th-century vintage as ‘scientists’ themselves, and intended to refer to the confluence of science and the ‘mechanical arts’ in awe-inspiring systems like the railroad and telegraph. But the term has a curious double valence. It seems to mean at times the machines themselves; on one level, this is how the press approaches the ‘hazards’ of technology. For the debates I consider, the focus is typically on risks posed to our minds by computers and televisions and increasingly ubiquitous electronic devices. I follow Marx, however, in emphasizing a different sort of hazard in the concept of technology. Your PC, just like a locomotive, would not function without a whole associated sociotechnical infrastructure: networks of cables, of manufacturing facilities and corporate financing, of trained engineers, competent users, and so forth. When we say that ‘technology is changing our

lives,' we're both right and wrong. Technology is not an autonomous force shaping our selves and societies from without, but neither is it a mere tool subservient to whatever ends we choose.

Technology is better understood as a name for social processes whereby humans and designed artifacts interact and reciprocally influence one another. Provided we keep this in mind, we are well positioned to understand the effects of media without succumbing to the hazardous form of technological determinism.

Mediation, ANT and STS.

This brings us to another relevant sense of 'mediation,' that proposed by Bruno Latour and Actor-Network Theory (ANT). I subscribe to a generally philosophical reading of Latour's project, in which despite all its shifts in emphasis and vocabulary, the essence of the theory remains an outgrowth of his epiphany as a youth: "I knew nothing, then, of what I am writing now but simply repeated to myself: 'Nothing can be reduced to anything else, nothing can be deduced from anything else, everything may be allied to everything else.' This was like an exorcism that defeated demons one by one" (Latour, 1993a, p. 163; Harman, 2009). Through his early collaborations with Michel Callon, whose notion of translation became central to ANT, as well as John Law and Steve Woolgar in the Anglophone world, this epiphany was developed through the 1980s and 1990s into a general descriptive theory which became a touchstone for the nascent field of science studies. Compared with the stable edifice of its chief antagonist within the field, namely David Bloor's strong programme, ANT is notable for its slipperiness.¹² The works inspired by ANT are united not by a

¹² In their well-known polemic the authors lay out the fundamental terms of the opposition – Bloor rests his case on a sharp divide between 'nature' and 'beliefs about nature,' which Latour sees as misleading – and equally trade barbs about Latour's tendency to change his mind about things, as opposed to Bloor's stubbornness (Bloor, 1999; Latour, 1999). Latour once famously said there were four things wrong with ANT – 'actor,' 'network,' 'theory,' and the hyphen (Latour,

specific method, fixed concepts, or even a style of explanation, but rather a skepticism toward explanations that invoke sharp ontological divides: for instance between human and nonhuman, natural and cultural, real and artificial. Instead the world of actors is flat, with all entities placed on an equal footing. There is no reduction or deduction, no one privileged set of things that explains the actions of all the others, and yet through a work variously defined as translation, *interessement*, and mediation, in principle any actor may mobilize a powerful network of other entities. Politicians, ideologies, minds, fictions, microbes or cosmic rays may all equally have effects, but only if we can describe the chain of translations by which diverse goals are allied together, and the mediations through which the actors on different scales communicate.¹³

As John Law puts it, “the stuff of the social is not simply human.” Instead, the social is conceived as “nothing other than patterned networks of heterogeneous materials . . . these networks are composed not only of people, but also of machines, animals, texts, money, architectures” (1992, p.2). This is the move for which ANT is best known: in its actor-networks, a wide array of biological, material, and machinic entities are granted agency on par with that of human beings. Such a move is of course deeply problematic for many of a humanistic bent, often provoking rebuttals that ‘only humans have intentionality!’ The response is twofold: in the first place, the agency of its

2007) – a claim which can be read charitably as caution against reifying any of its constructs, but equally one of many such remarks which incense those who seek to engage critically with the theory. Some of this vacillation, and the rather confrontational tone of Latour’s ‘introduction’ to a decades-old theory (Latour, 2007) are no doubt tied to his worries about how the uptake of ANT in the STS community has diverged from his original epiphany.

¹³ Graham Harman provides a useful illustration of these two terms by way of military tactics: “the means of linking one thing with another is *translation*. When Stalin and Zhukov order the encircling movement at Stalingrad, this is not a pure dictate trumpeted through space and transparently obeyed by the participant actors. Instead, a massive work of mediation occurs. Staff officers draw up detailed plans with large-scale maps that are then translated into individual platoon orders at the local level; officers then relay the orders, each making use of his own rhetorical style and personal rapport with the soldiers; finally, each individual soldier has to move his arms and legs independently to give final translation to the orders from above. Surprising obstacles arise, and some orders need to be improvised—the enemy melts away at unexpected points but puts up stubborn resistance in equally startling places . . . No layer of the world is a transparent intermediary, since each is a medium: or in Latour’s preferred term, a *mediator*” (Harman, 2009, p. 15). Mediation is therefore communication for Latour as well, but almost always with some degree of transformation, never a simple passing of messages.

actors is an entirely separate matter from intentionality, and second, this ‘intentionality’ is by no means the fundamental matter of fact which the humanists make it out to be.¹⁴ Thus ANT evidently “treads on a set of ethical, epistemological, and ontological toes” (ibid., p.3). In lieu of the distinctions long deployed by sociologists – for instance between natural, cultural, and technological phenomena, between human agency and its tools, between ‘macro’ and ‘micro’-phenomena, or even between ‘social causes,’ ‘social phenomena,’ and broadly non-social ones – ANT offers only this vision of a heterogeneous network, flattening out all such differences in kind. Latour insists that the social must be conceived as flat, without the usual sociological divide between micro-level interactions and the macro-scale ‘global context’ (2005, p.165). Even apparent unities and individual actors, whether ‘micro’ or ‘macro,’ technical, biological, or social, are to be conceived as punctualized networks. Televisions, human beings, and corporations can all appear as coherent unities, and we daily interact with them as such, yet in cases of breakdown it quickly becomes obvious that they are in fact networks of subcomponents. The TV needs a new circuit board; brain trauma leads one to forget their own identity; or perhaps the bank teller exploits the position granted them to steal your identity. In each case, we can’t help but recognize the complicated network behind what was once punctualized as a unified actor. This reveals the simultaneously social and technical character of entities which we sometimes gloss over as *simply* one or the other.

¹⁴ This is especially relevant in the context of cognitive science. See, for instance, Latour’s comments in *Reassembling the Social* on the idea of action being ‘overtaken’ (2005, p.44–62), where he points out that even the agency of a human is by no means reducible to intention or will, but on the contrary is itself produced by a network according to scientific accounts. The human, in this case, is simply a punctualized network: a network of brain, nerves, and muscles which, provided it functions properly, presents itself as a unified whole. “What counts as a person is an effect generated by a network of heterogeneous, interacting materials” (Law 1992, p.4; also Harman, 2009, p. 128): this illusion falls apart, of course, in cases of serious injury or trauma. See also (Dennett, 1989), where intentionality is described as an interpretive stance one can take *with respect to* a given system (human or otherwise), rather than an intrinsic or essential characteristic of such a system. I develop the connections between ANT and cognitive science in the embodied, extended paradigm further below.

As with McLuhan's probes, I draw out some relevant and useful elements of this account in developing my own project. I do not attempt to adopt or apply this theory as such – an enterprise which Latour himself cautions against (Latour, 2007), and a dubious one in any case given its polymorphous nature. What I take from this account is that the realm of the 'social' is not a distinct domain of human actions and intentions, but a field of interaction shaped equally by designed artifacts and other non-human actors. And the *way* that these elements interact is through mediation and translation. In principle this is a flat ontology which positions us on the same plane as machines and other nonhumans. Nonetheless, humans do have a privileged position as agents capable of recruiting many other actors to do our bidding. This isn't something in our nature, however, handed down as in Genesis, but a position which we must continually labour to maintain.

Turning from the general outlines of ANT to Latour's accounts of technological effects on human activity, they have less to do with the 'ratios' between different modes of sense perception, as in McLuhan, and more with how scientific laboratories have managed to extend our sensorium into a multitude of previously-inaccessible domains.¹⁵ He also draws on the psychologist J.J. Gibson's notion of 'affordance' to provide a fuller account of the processes by which technologies can affect what we do:

Technologies bombard human beings with a ceaseless offer of previously unheard-of positions – engagements, suggestions, allowances, interdictions, habits, positions, alienations, prescriptions, calculations, memories. Generalizing the notion of affordance, we could say that the quasi-subjects which we all are become such thanks to the quasi-objects which populate our universe with minor ghostly beings similar to us and whose programmes of action we may or may not adopt. (Latour, 2002)

¹⁵ Latour does address something akin to sensory 'ratios' in various ways throughout his work, most notably in ascribing a central role within science and technology to the training of senses to pick out increasingly fine distinctions: a matter of "bodies learning to be affected by hitherto unregistrable differences through the mediation of an artificially created set-up" (Latour, 2004a, p. 209).

Like McLuhan, Latour rejects a simple instrumentalism. Technologies are no ‘mere means,’ whose consequences are wholly determined by human goals and intentions. Rather, they often direct our intentions in various ways, and may be seen as ‘delegated’ actors with goals of their own—either incidentally, or having been explicitly designed in, as in Latour’s favourite example of the speed bump or ‘sleeping policeman’ (Latour as ‘Johnson,’ 1988). McLuhan likewise discusses these kinds of technological effects, but in terms of offering an explicit account of *how* media exerts its effects, focuses most on the somewhat inadequate theory of sense ratios, along with an equally dubious neurophysiological account.¹⁶

Technologies are still less ‘extensions’ of humanity for ANT, instead more likely to substitute *for* humans in contexts where our capacities are insufficient (durability, reliability, attention, perception, etc.).¹⁷ ANT is a crucial supplement to McLuhan in this regard, and further problematizes the opposition between technical and social determinisms. Latour sees a better account of technical mediation as the necessary solution to the ‘twin mistakes’ of both. Paraphrasing the political cliché, it is not the case that ‘guns kill people,’ nor that ‘people kill people,’ but instead we must consider what programmes of action are afforded to “the hybrid actor composed . . . of gun and gunman” (Latour, 1994, p. 33), as compared with other sorts of sociotechnical assemblages. Actors may be human, nonhuman, or hybrids of the two, and actions result from a play of resistances between the goals of each (or if we’d rather avoid ‘intentional’ language, their “functions, as engineers prefer to say:” *ibid.*).

¹⁶ I explore this aspect of McLuhan’s later writing further in subsequent chapters.

¹⁷ Thus, far from seeing ‘anthropomorphism’ as a mistake in understanding technology, when properly conceived, Latour sees the concept as perfectly appropriate. *Anthropos* and *morphos*, taken together, mean either what “has human shape or gives shape to humans,” and so he views even a simple technology like a door-closer as anthropomorphic in at least three senses: first, shaped in design and construction by humans; second, substituting for the actions of humans; and finally offering its own ‘prescriptions’ on what actions humans may take in relation to it (Latour as J. Johnson, 1988).

Though Latour entitles his introduction to actor-network theory ‘Reassembling the Social,’ his chief targets in the volume are those he labels “sociologists of the social:” theorists who define ‘the social’ as a transcendental domain or a ‘special kind of stuff,’ whether of explanations or *explananda*. Such theorists, he argues, tend to refer to such concepts as ‘social causes’ (as tacitly opposed to physical or psychological ones), ‘social factors’ or ‘social influences’ (as opposed to natural or biological ones), or ‘social structure’ (as opposed to material or technological: Latour 2005, p.3 *et passim*). Latour follows Gabriel Tarde and Harold Garfinkel in rejecting such oppositional views of the social with their implied nature/culture divide, proposing instead a ‘sociology of associations’ which defines the social as a “type of connection between things that are not themselves social” (2005, p.5). Hence ‘reassembling the social’ is not a project which Latour claims to have completed in his texts – ‘here, now, is the essence of The Social, which you can now apply as a framework to your study’ – but an ongoing labour. The task of the sociologist, as he frames it, is to trace the assembly of new associations from and by human and non-human actors. This continuing work of assembly (rather than the completed assemblages themselves) simply is the social, according to Latour and ANT.

The consequence is a radical extension and simultaneous limitation of the term. On the one hand, Latour notes that if ANT appears to ‘dilute’ sociology to the study of any “aggregate, from chemical bonds to legal ties, from atomic forces to social bodies,” this is “precisely the point that this alternative branch of social theory wishes to make” (2005, p.5), as a wide range of such aggregates may be assembled into a given actor-network. For him “an atom is no more real than Deutsche Bank or the 1976 Winter Olympics, even if one is likely to endure much longer than others” (Harman 2009, p.15). The very sparseness of the actor-network model leads to the affirmation of a vast range of phenomena as both ‘real’ and ‘social,’ and the simultaneous rejection

of many cherished dichotomies: human and non-human, of course, but also ‘subject’ and ‘object,’ micro and macro, ‘Nature’ and ‘Culture,’ and so on. Conversely, any invocation of ‘the social’ qua reified generality or ultimate source of causal effects is rejected out of hand by ANT.

While ANT’s applicability is not limited to science studies, it does have a special relationship with the sciences. According to Latour, it was the encounter with science that presented the *reductio ad absurdum* of traditional social critique: while it seemed perfectly reasonable to decry primitive divinities or regressive political movements as epiphenomenal projections of ‘social forces,’ he and the other theorists associated with ANT came to feel that it was useless to try and ‘crack open’ the sciences in similar fashion, to reveal “religion, power, discourse, hegemony” as causes of scientific knowledge. As he goes on to argue, “critique was useless against objects of some solidity” (Latour 2004, p.242). The ‘projection trick’ and the notion of ‘social cause’ simply don’t work on science, he claims, because scientists are constantly grappling with, even directly intervening in, the real material of the world. Accounts of reality are of course underdetermined by the data we obtain, but the concept of the real is defined here as that which resists our ability to form *just any* interpretation, and we come to know this reality by the mobilization of networks. The community of scientists plays a crucial role in this regard, but as they are always the first to emphasize, so too do the unpredictable responses of the agencies in their Petri dishes or brain scanners or computer simulations.

This is what revealed the poverty of the ‘critical trick,’ according to Latour, not only in the case of science studies, but of sociology and social theory more generally. Its tendency toward the debunking or ‘unmasking’ of specific claims or practices as the effects of ‘social causes’ is a principal target of ANT. When Latour claims, for instance, that “Laboratories are now powerful enough to define reality,” (1987, p.93), it is a serious mistake to read this as a slight against scientific realism, or an invocation of ‘social constructivism.’ On the contrary, his point about reality being “what resists”

(*ibid.*) is simply that laboratories, by constituting a set of alliances between humans and non-human actors, and by selectively manipulating the resistances of each, are now able to produce privileged determinations of the nature of reality. Thus, while he claims that “most philosophical discussions of realism” have been “very unrealistic” (1999, p.24), he can justly describe his approach as attempting to develop a more “realistic realism” (*ibid.*, p.15).

Media technologies thereby play a crucial role in society as conceived by actor-network theory. From its earliest days it has accorded a central role to *inscription devices*, and to the production of printed publications (Latour, 1979). The process of producing scientific knowledge is stripped down in early ANT accounts to a chain of translations or mediations, whereby the traces of some invisible agency are made visible through technologies, organized into scientific publications allying textual argumentation with graphical demonstration, and ultimately then translated into the kinds of factual claims presented as settled in textbooks. Communications media, especially print and graphical inscription, are thereby as central to the ANT account of science as they are for Eisenstein and McLuhan.¹⁸

ANT also proposes a special sense of the term mediator, however, as opposed to ‘intermediary.’ This designates a category broader than the typical sense of media, of actors which transform the flows passing through them rather than serving as mere passive conduits.¹⁹ Summing

¹⁸ The latter contends for instance that “the art of making pictorial statements in a precise and repeatable form is one that we have long taken for granted in the West. But it is usually forgotten that without prints and blueprints, without maps and geometry, the world of modern sciences and technologies would hardly exist” (McLuhan, 1994, p. 157). Here, McLuhan seems to closely prefigure Latour’s “Visualization and Cognition” (Latour, 1986) and a whole ensuing tradition within STS.

¹⁹ Actors of the latter type are labeled ‘intermediaries,’ and in Latour’s view they are relatively rare – transformative mediators are more typical, and where intermediaries are in operation this is not a matter of their pre-given ‘essence’ but an ongoing labour by other actors. An oil pipeline, for instance, is a paradigmatic ‘passive conduit,’ but keeping it running as an intermediary is a costly process involving physical inspections, monitoring devices, specially-designed coatings and cleaning devices, each depending on a multiplicity of subsidiary actors. This upkeep process can of course fail, and the pipeline will then end up transforming its flow in a catastrophically unplanned and undesirable fashion.

up the value of ANT to my project, the constitution of an actor-network to produce claims *about* media effects is best seen as a process where certain nonhuman actors (brains, brain images, genes, televisions, video games, etc.) and human actors (basic researchers, clinicians, behaviourally disruptive children, journalists, etc.) are assembled together to construct different accounts of reality: sometimes serving as passive intermediaries in the stabilization of one, sometimes reshaping another or presenting their own resistances, with roles shifting over time. Macro-scale forces like ‘corporate power,’ ‘profit motive,’ ‘moral responsibility,’ ‘innate temperament’ and so forth are not suitable explanations for the entire process, but various actors may attempt to recruit them into their accounts of the world, with varying degrees of plausibility.

This is a more challenging picture of knowledge and truth than one limited to a humanistic sociology of beliefs and interest groups, but in my view it does greater justice to the ways that technologies shape society. In reading Latour alongside McLuhan, it seems they are working toward the same fundamental problem from very different starting points. McLuhan comes from the world of literary studies, where a specific mode of rationality derived from print is ‘naturalized,’ so to speak; Latour from sociology in the tradition of Durkheim, where the very notion of the social is reified and naturalized. But both work against these traditions in a drive to analyse how media – ‘inscription devices’ and beyond – act in a structuring, performative fashion, shaping our collective possibilities for cognition and (inter)action. The point is precisely not that technology determines the course of human society, but that communications media in the broadest sense are the technical supports by which a distinctly human society is assembled.²⁰ To put it in something more like

²⁰ And as Latour has suggested in his work with the primatologist Shirley Strum, we should not view sociotechnical interaction as exclusive to human collectives, but rather we may observe a gradation of complexity in the materials employed to stabilize the socius, from bodies alone (in bees and baboons alike) to more complex tools (in chimpanzees: (Strum & Latour, 1987).

McLuhan's probing style, we might say that human society is itself the *arché*-effect of media, while media are equally designed, maintained, and transformed through social interaction.

There is one final relevant sense of the term medium which again differs slightly from and overlaps with the common, the McLuhanesque, and the ANT senses: that derived from computer science and information theory, where a medium is any physical carrier of an encoded form of information. This is the sense in which 'television' is not a medium but 'co-axial cable' or 'radio-frequency waves' are, 'print' is not and 'parchment' or 'papyrus' are. McLuhan decried the Shannon-Weaver theory of communication as "an extreme example of the lineal bias in communication," and of the Western print-derived privileging of 'software content' over 'hardware container' (McLuhan, 1978, p. 58), which even as a 'probe' seems to fundamentally miss the point of the model's particular utility within engineering.²¹

It is this longstanding sense of a medium as a vehicle or substrate which is most central to my claim in the next section that cognitive science articulates a distinctively *mediated* understanding of cognition. The founding assumption of orthodox cognitivism is that 'mind' itself may be a matter of software content, separable from its familiar biological hardware and susceptible to at least modeling, if not replication, in electronic computers. Mind thus becomes a kind of signal which can be realized in multiple distinct media. Conceptual metaphors of information and encoding are as central in the development of cognitive science as in molecular biology, in both cases following from the ideas of the cybernetics group while adapting them in distinct ways (Miller, 1953; Dupuy, 2000; Kay, 2000). In that same essay, McLuhan evidently saw the merits of an engagement between media theory and neuroscience, yet unfortunately the one he proposed was founded on a relatively

²¹ This theory returns as a point of reference for actors throughout this text, and I address its implications for cognitive science more fully in the next chapter.

immature and problematic conception of hemispherical localization, and a dubious cultural binary: the West is ‘left-brained,’ the East is ‘right-brained’ (McLuhan, 1978, pp. 54–55). While I aim to show the development of cognitive science is patterned by mediations on a number of levels, I also intend to employ its findings in a more sophisticated way to supplement my account of what effects the media may have.

This project is thereby constituted according to a double circularity. In the first place, I want to extend the project of Innis and McLuhan, which I have argued is closely aligned with that of Latour: to examine the consequences of technical mediation for mind and society, in the era of the Internet and other increasingly ubiquitous modes of digital communication. Vast in scope, this is equally a problem for media studies and for science and technology studies. Scientific research is positioned as a privileged discourse of truth about minds in our society, and nonetheless directly subject to the same kinds of media effects. Hence I also consider these effects within the practice of science, at once informed by and feeding back into my analysis of media. These circular processes delimit and focus this project, as I examine the mediations involved across the history of cognitive science, its own investigations of the effects of media, and ultimately the dissemination of its theories, metaphors, and images within the public sphere. In the end, I do not dispute that research within cognitive science is among the most effective ways to address the question of media effects. Carefully examining the effects of media within the practices of science should not undermine but augment the plausibility of such research, perhaps suggesting new directions for future investigations or new interpretations of old ones (Choudhury, Nagel, & Slaby, 2009). To that end, I wish to incorporate elements of cognitive-scientific accounts within my own analytical vocabulary.

Hence I seek to execute what has become a fraught maneuver for scholars of science and technology: to interrogate and problematize a strand of scientific research, while simultaneously

drawing upon its findings and its ongoing debates to inform my conclusions. I will at times be unpacking and redeploying positions and critiques internal to the field, rather than limiting myself to an external, sociological meta-discourse. Drawing on scientific accounts in this way has recently been criticized as a dangerous regression by Trevor Pinch, in commenting on Karen Barad's 'agential realism' (Pinch, 2011). The main ideas I employ in this fashion are present equally in the work of Andy Clark and like-minded cognitive scientists as in the writings of McLuhan and Latour, namely that human cognition emerges through and is continually shaped by collective interactions with technical media (Clark, 2003). In that respect as well it is like the quantum mechanics which Barad draws upon in devising her theory: the data of both fields are only comprehensible with a considerable supplement of philosophical interpretation.²²

Without wishing to adopt Barad's broader theory here, I concur in her rejection of Pinch's categorical claim of 'mutual exclusivity' between the results of science and the study of scientific practice (Barad, 2011). My approach is less directly tied to a particular account (like Bohr's philosophy-physics for Barad) and instead adheres to basic principles inherent in ANT and ethnomethodology: starting out from some of the ways in which various cognitive scientists and other relevant actors have described the interactions between mind and media, and seeing what these can reveal about research practices as well as processes of dissemination. Fundamentally, moreover, my characterization of these processes is closely aligned with the core ideas of cybernetics, as enshrined in the name of its Macy Conferences—they are a matter of 'circular causal systems,' cycles of multidirectional feedback and reciprocal interaction. Pinch's categorical

²² Leon Eisenberg also presents a "more suggestive than precise" parallel with Bohr's complementarity principle in the context of psychology, arguing that the 'mind' and 'brain,' while seemingly contradictory explanatory schemas, are both as necessary to account for psychological phenomena as the wave and particle theories are for micro-scale physical phenomena (Eisenberg, 1986). Notions of 'intra-action' and entanglement are in my view well suited to describe both sets of phenomena, but this remains just a suggestive analogy, given that despite much speculation no convincing evidence has been presented thus far for quantum mechanics playing any role in brain function.

bracketing seems unnecessarily limiting for STS; not all circularities are dangerous. If science suggests that media have certain effects on cognition, then why should we not reflexively consider what implications this may have for the profoundly mediated cognition of scientists?²³

An understanding of cognition as embodied, embedded, and distributed provides a way to resolve a number of tensions in this regard. I am by no means the first to propose an alliance between this paradigm and the study of science and technology. Actor-network theory has long recognized this possibility, after an initial skepticism. Latour and Woolgar once famously called for a “ten-year moratorium on cognitive explanations of science,” promising that “if anything remains to be explained at the end of this period, we too will turn to the mind!” (Latour, 1979, p. 280). With ‘cognitive explanations’ construed according to the ‘heroic genius’ school of history and sociology of science, this caution was no doubt a valuable one. I have no interest in cognitive accounts taking this form, nor in adopting what amounts to a more advanced variant of this idea with the cognitivists’ technological metaphor of mind as information-processing in the brain. Instead, I follow Latour’s past-due recanting of the ‘moratorium’ in his review of Hutchins’ *Cognition in the Wild*: with the contention that cognition itself is a mediated phenomenon, reliant on the “propagation of organized functional properties across a set of malleable media” (Hutchins, 1996, p. 312), the great divide between mind and society disappears. Cognitive explanations in this distributed, embedded form have thereby been “made thoroughly compatible with the social explanations of science, technology and formalism devised by my colleagues and myself” (Keller, Bazerman, & Latour, 1996, p. 62) –

²³ In contrast with Barad, I am not trained nor do I position myself as a cognitive scientist, and the expertise I cultivate is appropriately described in Harry Collins’ terms as ‘interactional’ rather than ‘contributory.’ It is worth noting, however, that cognitive scientists themselves exhibit no concern over conducting this kind of circular analysis, and a substantial body of work is developing in the ‘cognitive science of science’ (Thagard, 2012). Jamie Cohen-Cole also outlines a similar view: “work in science studies has argued for analyzing science by breaking down the analytic boundaries between the natural world and the social world (Latour, 1993). If applied to history of psychology, this work would imply the analytic value of breaking down distinctions between the natural world described by psychology (the human mind) and the minds and social worlds of psychologists themselves” (Cohen-Cole, 2005, p. 109).

though eventually Latour will set aside the language of ‘explanation’ altogether for its unfortunate ‘debunking’ connotations as described above.²⁴ Be they explanatory or descriptive, on this view there need be no opposition between cognitive and social accounts of science, and I contend that this is a sufficiently plausible and fruitful working theory for general adoption within STS. I employ it, however, in its broader and more philosophical form, as the idea that minds, societies, and technologies stand in an ongoing emergent relation of co-construction, not fully ‘determined’ by intrinsic properties of any one domain. This does not stand or fall, as Barad’s approach seemingly does, on the results of specific scientific controversies, for instance over how ‘radically’ embodied and distributed cognition really is (Adams & Aizawa, 2008; Chemero, 2009). I embrace the fairly uncontroversial thesis that our minds and bodies are naturally primed to extend and distribute our capacities through tool-use: an innate tendency toward hybridization which gives rise to sociotechnical cognition and culture.

The sciences and economies of attention.

McLuhan once mused in a letter that his ‘style’ was widely misunderstood, but nevertheless it was a “very good style for getting attention” (Molinaro, McLuhan, & Toye, 1987, p. 505). Attention itself is another concept which unifies and delimits this project. In the final section, I conduct a case study of the medicalization of attention: again focusing on a recursive interaction, I examine how media stories comment on the potential causal role of media use in the development of attention

²⁴ Theodore Porter has recently argued on the basis of this comment that both Latour and Woolgar’s early methodology and the ‘often motiveless and detached from material and social circumstances’ cognitive explanations stand as models of troublingly ‘thin’ descriptions. Though valid as a critique of the worst tendencies in both actor-network theory and cognitive science, Porter unfortunately ignores this recanting and *rapprochement* between cognitive, social, and material accounts, which in my view already offers a deeper argument for the ‘full compatibility’ between intellectual history and ‘an interest in objects and practices’ which he proposes (Porter, 2012).

deficit disorder (ADD). But even considering the consequences of media for ‘normal’ cognition, our questions always circulate alongside questions of attention. How do the multitude of devices to which we pay attention act upon us? More directly, what does it mean that the citizens of advanced industrial democracies are surrounded by so many spectacular ‘attention-grabbing’ phenomena in our everyday lives? This is evidently a major popular concern, as evidenced by media accounts I examine throughout this project. In a way, addressing these questions is also my aim here. Yet in so doing I wish to critically interrogate the popular framings of these issues, with particular attention to the ways in which they invoke neuroscience and cognitive theories.

As Jonathan Crary suggests, “Attention is not just one of the many topics examined experimentally by late nineteenth-century psychology but is the fundamental condition of its knowledge” (Crary, 1999, p. 24). Throughout their history, research techniques in the sciences of mind have been grounded in attention more than any other cognitive process. The early, paradigmatic investigations of reaction times, sensation and perception, reflex, and conditioning all “presupposed a subject whose attentiveness was the site of observation, classification, and measurement, and thus the point around which knowledge of many kinds was accumulated” (*ibid.*). This continues to the present day, as nearly all research into ‘normal’ brains measures neural correlates of subjects’ attention to some particular stimulus, often presented as well as measured through digital media. Simultaneously, one of the most widespread forms of mental pathology is said to be a deficit of attention. My discussion is informed by Crary’s analysis, though my focus is on more contemporary mediations of attention. I likewise understand the problem of attention as elaborated within a specific socio-economic regime, one that demands “attentiveness of a subject in a wide range of new productive and spectacular tasks, but whose internal movement was continually eroding the basis of any disciplinary attentiveness” (Crary, 1999, p. 29). Adherence to the demands

of this regime on the sides of both production and consumption demands that “we accept as *natural* switching our attention rapidly from one thing to another,” in an ongoing play of “reciprocal attentiveness and distraction” (*ibid.*, p.30).

In its conceptualization of both normal and pathological minds, cognitive science is enmeshed within these broader political-economic and sociotechnical fields, and can never be truly purified of such associations. The brains and minds it studies are inevitably shaped by interactions with a range of media, particularly given conjoined notions of neuroplasticity – that is to say neural growth and ‘rewiring’ in the adult brain – and extended cognition. The pathologies which it diagnoses are inevitably linked up to failures of adaptation within an economic system now fully oriented toward attention and explicitly theorized as such by prominent actors within it (E. Dyson, 2012). In the case of ADD/ADHD which I take up in greater detail below, this is often characterized by an excess of capture on the consumptive side of an ‘attention economy’ – by children ‘fixated on screens, and nothing else’ (Klass, 2011) – implying a concomitant failure to meet the demands of productive labour. What perhaps characterizes the contemporary situation above all, though, is that this supplanting of productive energies by consumption to pathological excess does not present a ‘social problem’ in the same sense as it once did. The industries of attention capture are so economically central and well-positioned to leverage this consumption behaviour that business now depends as much on the leisure of the masses as their labour.

Bernard Stiegler has recently written an extended essay on the political stakes of attention, arguing against what he perceives as an ongoing destruction of our attentional capacities by the culture industry. He deems this a ‘war for intelligence’ and a catastrophic threat to the rational traditions of the Enlightenment - a tradition which, for all its deconstructions within the French tradition Stiegler represents, he insists remains worth defending (Stiegler, 2010). This inverts the

‘attention economy’ hypothesis, proposing that as market actors develop increasingly sophisticated techniques for capturing attention, it represents a mounting threat to those capacities for ‘mature’ attention (and, by extension, cognition) which have been recognized since Kant as central to democratic self-governance.²⁵ Stiegler’s analysis transposes the Derridean notion of the *pharmakon* into the contemporary media context, arguing that technologies of communication are simultaneously cure and poison for the human mind, at once fundamental to its ‘individuation’ and development but also profoundly destabilizing. Equally, he reframes an omnipresent moral panic – that new media are harming children – within the strange vocabulary of French philosophy. His claim is not simply that culture industries are harming children, but rather they are making children of us all, leaving us collectively unable to ‘emerge from our immaturity’ and to constitute an enlightened, self-governing *polis* according to the Kantian ideal. “Taking care of youth and the generations,” the title of his essay, is a task which demands responding to this challenge, with ‘care’ in the Heideggerian sense [*Sorge*] understood as a matter of shared responsibility to posterity, an ‘infinity of generations’ (Stiegler, 2010, p. 186). The challenge posed by an economy of attention is how to counteract the ‘psychotechnologies’ of attention-capture and redesign the institutions by which we transmit across generations the habits of sustained, rational attention characterizing a literate public sphere. I concur with Stiegler that this is the fundamental challenge facing our educational and psychiatric institutions today, and that their responses – typified in the former case

²⁵ Here Stiegler is referring to the notion of enlightenment as ‘emergence from self-imposed immaturity,’ and the ensuing constitution of a collective capable of ruling itself. Like McLuhan, Stiegler is explicit that these practices of rationality are specifically grounded in the technology of print, reproaching Foucault for neglecting this dimension of Kant in his own ‘What is Enlightenment?’ this maturity is that of “a mature consciousness *that writes* before a public of mature consciousnesses *able to read those writings*. This maturity is technological because it is inscribed within an *apparatus of writing* that is also a *society*” (Stiegler, 2010, p. 116). Though his framing is complex, a brand of media theory derived from Derrida, his critique here recalls any number of comparable *cris de coeur* (since at least the Frankfurt School) for the decline of a rational reading public brought about by media industries.

by dubious neoliberal ‘reforms,’ and in the latter by the diagnosis of attention deficit disorder – have been woefully insufficient thus far.

Stiegler does not fall into the same traps as many otherwise similar conservative laments for the displacement of one medium by another. It is not that the advent of television and the Internet corrupted a previously natural, authentic, and unmediated mode of attention and cognition; rather, like McLuhan, he recognizes that habits of attention have always emerged from a “process of psychic individuation, collective and technical” (Stiegler, 2010, p. 29). The rationality of the Enlightenment and of liberalism is not a transcendent absolute but a contingent product of print culture. Perhaps his most valuable insights are encapsulated in three points which, toward the end of the text, he emphasizes as the prerequisite to any thinking about ‘responsibility’ in our era:

1. That what Heidegger calls *Sorge*, taking charge of one’s existence as one’s own...is constituted as attention in that attention is always already at once psychic and collective...
2. That attention is thus precisely *constructed technically* meaning that any ethics is in an essential relation to technics...
3. That attention emerges from a formational process that is a social organization, constitutive of the transindividual and transindividuation, transindividuation being transmitted as much technically as ethically from generation to generation (Stiegler, 2010, p. 186).

Throughout Stiegler’s philosophy there runs the same current as we find in McLuhan and Latour, of concern with the originary and continual shaping of the psychic and the social through media (or ‘technics’). He emphasizes, moreover, how this shaping is at once an individual matter, a process of ‘individuation’ or subject-formation, and equally a social matter, of *transindividuation* through interaction across space and time. The real ethical concern is not whether new technologies are corrupting some mythically ‘authentic’ human experience, for cognition is essentially cultural and constructed. Instead we must examine the distinct affordances of old and new technologies, and design educational and social systems such that the worst tendencies of new media forms are mitigated, the best of old forms are preserved, and the youth are not rendered chronically,

pathologically deficient in their capacities for sustained attention. Invoking both the Platonic-Derridean *pharmakon* and contemporary medical discourse, Stiegler described this project as a ‘pharmacology’ of mediated attention. While new technologies as deployed by commercial interests seem to be dramatically supplanting human attentional capacities in a variety of ways, he sees them as equally containing the conditions of possibility for vital new forms of social intelligence, collective cognition, and political action.

In all these respects, Stiegler’s consideration of the contemporary political economies of attention is laudable, and informs my own analysis. In other respects, however, his framing is more dubious. I am less convinced that it makes sense to view our world in terms of a grand ‘battle for intelligence,’ of a liberatory Enlightenment rationality against the pernicious influences of a metastasized culture industry. Stiegler’s analysis is informed by Katherine Hayles’ conceptualization of ‘hyper’ and ‘deep’ attention as a ‘generational divide’ between two distinct cognitive styles: the latter being more typical of the sustained attentiveness demanded by print culture, while the former constitutes a rapidly-shifting vigilance over multiple competing foci of attention, more associated with digital culture (Hayles, 2007). Despite all the nuances of his critique, Stiegler seems insufficiently open to the value of ‘hyperattentiveness’ within our culture, or to generational shifts in patterns of attention, away from the model which privileges sustained, deep attention above all, toward a cognitive style which is less familiar but may be nevertheless fruitful in terms of human autonomy and cultural flourishing.²⁶ No doubt it would represent a substantial loss if the modes of

²⁶ Hayles describes the two modes through the vignette of a college sophomore immersed in *Pride and Prejudice* sitting on the couch alongside her younger sibling mashing buttons in *Grand Theft Auto* – the former obviously typifying deep attention, and the latter ‘hyper.’ She suggests herself that each is suited to different tasks: “Deep attention is superb for solving complex problems represented in a single medium, but it comes at the price of environmental alertness and flexibility of response. Hyper attention excels at negotiating rapidly changing environments in which multiple foci compete for attention; its disadvantage is impatience with focusing for long periods on a noninteractive object such as a Victorian novel or complicated math problem” (Hayles, 2007, p. 188). Hyperattention, she suggests, is also the more evolutionarily primary mode. This is the same sort of ‘attention’ an animal pays to its surroundings, continually alert to

reasoning derived from Enlightenment print culture were lost altogether. But just as printed words gained new vitality as hypertext on the Internet after their seeming decline in the era of television, it seems likely that deep attention will endure in new hybrid forms (perhaps, for instance, with the aid of simple artificial cognitive agents, as in the data-mining approach).

Stiegler is also unfortunately aligned with a certain French philosophical tradition that emphasizes creative redeployment – and occasionally sensationalistic reinterpretation – of scientific findings over critical analysis of scientific discourse. As Latour and Geoff Bowker pointed out reviewing the social studies of science in France, we may be used to seeing French intellectuals “appearing in the wings (when they do appear) as radical troublemakers of existing order,” and yet often “when science appears, these iconoclastic figureheads gather behind its banner” (Bowker & Latour, 1987, pp. 716–717). To be sure, there are exceptions to the rule, Latour himself being the most notable. But the principal issue with Stiegler in this regard is his uncritical usage of a study which I discuss at greater length later on in this project (Christakis, Zimmerman, DiGiuseppe, & McCarty, 2004). He cites this epidemiological study based on survey data as proving decisively that televisual media provoke “the literal destruction of children’s affective and intellectual capacities,” producing “dramatic increases in attention deficit disorder through the premature structuring and irreversible modelling of their synaptogenetic circuits” (Stiegler, 2010, p. 56). While his reading of the Christakis study is no doubt derived in part from overly credulous media reports of the type I analyse below, his notions of ‘synaptogenetic circuitry’ are drawn from research on neuroplasticity as interpreted by Katherine Hayles (Hayles, 2007). The study itself is survey research, having no direct connection at all to neuroscience.

potential danger; deep attention, Hayles and Stiegler agree, is the product of culture and educational institutions. Gilles Deleuze, incidentally, was fascinated by this former, primordial mode of attention, which he described with the phrase ‘*être aux aguets*,’ and which is closely tied with his conceptualization of affect.

Hayles' arguments about distinctive modes of 'deep' and 'hyper' attention are likewise valuable insights, and while also invoking scientific concepts in a relatively loose way, she does not draw overblown conclusions from meagre data in the same fashion as Stiegler. In part I simply don't share the apocalyptic outlook of Stiegler's narrative – I side more with McLuhan or Hayles in regarding the shifting forms of attention produced by new media as legitimate and valuable cognitive strategies in their own right, representing real gains as well as losses. Stiegler is focused more on the 'loss' column, and though he does so in a more sophisticated way than most, he still tends to mistake the displacement of a linear, print-derived rationality for the end of reason and attention as such. More importantly, however, I want to more carefully examine the discourse, practice, and mediated processes by which scientific claims like Christakis' are produced and circulated. These claims in themselves are controversial and not to be taken at face value, far less to be adapted into grand philosophical claims on the basis of simplified journalistic accounts.

Michel Foucault is another important influence, both for Stiegler and for this project. I return to his arguments most directly in the fifth section, when addressing psychiatry and shifting conceptions of mental pathology. His search for a 'common matrix' linking the 'history of the human sciences' with the 'more general field of other ways of exercising power,' be they punitive, political, governmental, or otherwise (Foucault, 1995), influences my analysis in a more fundamental but indirect fashion. Stiegler contends, however, that Foucault is not sufficiently attentive to the actual media technologies which underpin institutional archives, and that despite the value of his analysis of disciplinary society, he either never grasped or was not interested in the more recent shift to what Gilles Deleuze characterized as 'societies of control' (Deleuze, 1992). Disciplinary societies rested upon the partitioned ordering of bodies in space, the setting-up of ranks, functional sites, regimes of visibility and examination, as in Bentham's paradigmatic Panopticon. In the society of

control, apparently ‘freer’ movements of bodies are tolerated, on the basis of an ever more fine-grained and automated system of monitoring and governance: “The numerical language of control is made of codes that mark access to information, or reject it. We no longer find ourselves dealing with the mass/individual pair. Individuals have become ‘*dividuals*,’ and masses, samples, data, markets, or ‘*banks*’” (*ibid.*, p.5, emphasis original).

The society of control is intimately tied to the sciences of control – that is to say, what started off as cybernetics – and to the techniques they developed for transmitting, storing, and theorizing information in electronic media. The rise of control societies also runs in parallel with a shift in many parts of the world from an economy founded on and concerned primarily with exertion (Rabinbach, 1992), to an economy driven by attention, speculation, and intangible production. As Deleuze puts it, somewhat hyperbolically but perceptively, “capitalism is no longer involved in production, which it often relegates to the Third World ... What it wants to sell is services and what it wants to buy is stocks...marketing has become the center or the ‘soul’ of the corporation” (Deleuze, 1992, p.6). Later I discuss some of the ‘higher’ aspirations of Google, which perhaps more than any other corporation is emblematic of our present economy. But despite its dreams of artificial intelligence, it is always worth recalling that its most profitable line of business remains as a service provider for plain old marketing on the Internet. As the contemporary adage goes, ‘if the product is free, you are the product:’ their Web search makes them no money at all, except indirectly by supporting its AdWords program in offering some of the most effective tools for capturing, measuring, and monetizing the attention of a global user base potentially numbering in the millions.

This state of affairs provides the context for and grants urgency to the questions I ask here regarding media effects. We are confronted by competing and cooperating demands of productivity

and consumption, education and entertainment, reaching us through a plethora of networked screens. The boundaries between all of these domains are blurring. Psychology and the cognitive sciences link up with these commercial interests in a variety of interesting and troubling ways, forming a ‘common matrix’ of political-economic and rational-scientific apparatus (*dispositifs*, in Foucault’s terminology). They may be recruited within a society of control, where attention is among the foremost objects of monitoring, administration, ‘nudging,’ leveraging, and pharmaceutical treatment. Stiegler is not the only recent theorist to question this *dispositif* and draw strong prescriptive conclusions (Boltanski & Chiapello, 2007; Malabou, 2008). These critiques form valuable background to my analysis, but I aim to examine the interaction between cognitive science and media more closely and descriptively. I focus on the ‘media effect’ as public matter of concern, but perhaps also an ‘epistemic thing’ or an unstable ‘boundary object’ as constructed by scientific research (Rheinberger, 1997; Star & Griesemer, 1989). Whatever the terminology, this problematic entity constitutes the principal object of my study: what are the effects of media, and how do we come to know them? Particular devices do not have deterministic effects. Instead they act upon us within this common matrix of intellectual, economic, and political power, reorganizing our cognition and interaction in accordance with various goals. Sometimes we set these ourselves, sometimes they are set by various other actors with or without our knowledge and consent. The effects of technology are real; half the error of technological determinism is to misunderstand ‘technology’ as signifying mere machines, when really it should imply a whole sociotechnical *dispositif*. The other half, of course, is that far from being deterministic, these effects are often actively resisted.

In the next chapter, I present a historically-informed analysis of the role played by technological metaphor in understanding the mind. There I seek to explore how the media of communication have often furnished analogical material for our conceptions of mind and

subjectivity. The subsequent multi-chapter section considers the history of cognitive science as that of a profoundly *mediated* discipline, shaped by a long line of technological metaphors for cognition and technologies designed to implement those cognitive metaphors. Its three chapters respectively cover the eras of cybernetics, that of orthodox ‘cognitivism’ which literally equated cognition with computation, and the more recent challenges to that orthodoxy. The themes of attention remain prominent throughout. Developing in parallel with the shift from an economy founded on industrial production to an attention economy, I consider various aspects of the shift from a society centred on the ‘human motor’ to that of the ‘human information processor’ (Heyck, 2014; Rabinbach, 1992). The final sections return more directly to questions related to the popular media technologies competing for shares of our attention. The fourth chapter looks at what scientific research specifically addresses the effects of media on ‘normal’ minds, and how the effects of media are subtly at play in many other aspects of cognitive science. The fifth then considers the relationship between media and one of the most prominent mental pathologies of our time, attention deficit disorder. This section is divided into two chapters: the first covering the history of the disorder, of how attention came to be medicalized in this fashion, and the second undertaking a content analysis of print media bearing on the specific question of whether new technologies are believed to play a role in the etiology of this disorder.

2

Technological metaphor and the figuration of mind

“If the seventeenth and the early eighteenth centuries are the age of clocks, the later eighteenth and nineteenth centuries constitute the age of steam engines, then the present time is the age of communication and control.”

— Norbert Wiener (1961, p. 39)

In a historical overview of artificial intelligence research that he describes as more “grist for the historian’s mill” than a fleshed-out history in itself, Allan Newell notes that neither theories nor research methodologies are particularly satisfying ways to structure a historical account of his field. In this discipline which he pioneered along with contemporaries like Herbert Simon and John McCarthy, the theories put forth, especially when successful, were often inextricably “embedded in computer systems,” and he felt that in such cases the systems – and by implication what they could *actually do* – always seemed to “speak louder than the commentary” (Newell, 1982, p. 2). ‘Paradigms’ or ‘research programmes’ were concepts too coarse-grained for framing the history of AI, he argued, and the field has developed and maintained a single paradigm, or at most two.¹ Instead, we might better track the “history of AI as a whole... in terms of the geography of tasks successfully

¹ Here he was referring specifically to the split between symbolic, representational AI of the type he and Simon originally pursued, and alternative, connectionist, previously cybernetic conceptions of cognition – a divide which I cover in greater detail below. Yet he also rhetorically implies that they are perhaps not quite distinct paradigms, given that both have broadly conceived cognition in physicalistic, mechanistic terms, and share a number of core researchers and concepts.

performed by AI systems” (*ibid.*, p.6). In this chapter I begin to explore what such a history might look like, highlighting the crucial importance of computational tools for theorization in cognitive science. This is what the human actors throughout its history have always tried to affirm: that their ideas were constantly being guided and shaped by interaction with nonhumans, with their programs and their robots, as they have either cooperated with or resisted their theories of mind.²

Against this premise that ‘systems speak louder than commentary,’ however, I wish to juxtapose the potency of metaphor, a feature which some might consign to the realm of mere commentary. As much as it is a history of system-building, the history of cognitive science is the history of one generative metaphor, of mind as computation. Across this history the sense of the metaphor has shifted, as has that of computation itself. But the case always remains, as Kurt Danziger contends, that “the identifying relationships that occur in psychological discourse are of two kinds: There are explicit literal definitions, and there are relationships of metaphorical analogy. On the whole, the latter tend to be more pervasive than the former” (Danziger, 1990, p. 337). After examining some other takes on metaphor in science and psychology more specifically, I devote much of this chapter to some specific examples of technological metaphors in the discourse of mind which predate cognitive science proper: first from an era when computation was something done by specially trained humans, and then when the most fertile ground for analogy-making was the new domain of telecommunications. My analysis here is informed by the work of Michael Arbib and Mary Hesse, the most persuasive exponents of the view that “scientific revolutions are, in fact, metaphorical revolutions” (Arbib & Hesse, 1986, p. 156). Another touchstone is David Edge’s

² Already it should be noted that this aspect of cognitive science undermines the conventional split between ‘natural’ and ‘constructed,’ since what these tools reveal is only indirectly a matter of nature – and only if the analogy of mind with machine can be made to stick. The reality which cognitive science practically and directly reveals is a wholly constructed one, an abstract and functional description of the cognitive as instantiated in artificial systems. Yet as Newell’s longtime collaborator Herbert Simon and others have argued, this may be able to tell us a great deal about the architecture of natural minds (Simon, 1996b).

argument that in scientific research, “the successful metaphor does not merely provide answers to pre-existing questions: rather, by radically restructuring our perception of the situation, *it creates new questions, and in so doing, largely determines the nature of the answers*” (Edge, 1974, p. 136). As he goes on to note, such metaphors often derive from the ‘hardware’ of ‘control technology’ – this is particularly true in the cases of cybernetics and cognitive science.

Equally, however, Edge’s framing highlights some of the conceptual confusion which can result from an emphasis on metaphor. This is no minor problem, in that statements like the above often raise the ire of scientists, who again wish to privilege instead the observable actions of their nonhumans: for such minds, science is not about the superficial rhetoric one uses to promote or explain one’s work, but about what one’s programs can do. I draw upon some other influential past accounts of metaphor to argue in this section that – as in many other, if not all scientific fields – the role of metaphor in cognitive science is far from superficial. Metaphors have a ‘generativity,’ and are in some sense *constitutive* of theory (Danziger, 1990). They precede and shape the development of formal models. They allow for communication across scientific fields, particularly relevant in the development of ‘interdisciplines’ like cognitive science. Yet they also leave open precisely which elements of the source domain should be mapped on to which of the target. While they clearly do structure scientific questions, as Edge suggests, they do not tell us everything about scientific practice, and it is troublingly ambiguous to simply state that they “determine the nature of the answers” in science.

In discussions of metaphor in the cognitive sciences, we confront two quasi-misunderstandings, both with some merit and certain inadequacies:

1. “There is no reason why computation ought to be treated as merely a metaphor for cognition, as opposed to a hypothesis about the literal nature of cognition” (Pylyshyn, 1980, p. 114).

2. “Metaphors do not ‘express’ scientists’ ideas; they are the ideas. Metaphors suggest new visions, images, and models; they inspire scientists to approach problems in new ways” (Otis, 2001, p. 59).

It does little good to deny or suppress the generative role of metaphor in cognitive science, and it is simply false to imply that researchers have developed some complete ‘literal’ mapping of human cognitive processes onto artificial, computational ones. Despite some local successes replicating cognition-like phenomena with computer programs, the idea that *mind is computation* remains in many ways just a suggestive (and contentious) metaphor. Equally, however, it is a dubious rhetorical gesture to suggest that ‘scientists’ ideas’ are wholly metaphorical, or that metaphors shape scientific answers in a deterministic fashion. The metaphor of cognition-as-computation is indeed a *working* hypothesis within cognitive science, one which researchers are labouring on an ongoing basis to translate from the metaphorical and abstract into the literal and concrete. The success of this labour is by no means assured. The history of the field though is a story of successive groups each recruiting some significant actors – be they politicians, funding agencies, journalists or computational systems – to their particular vision of how this literalization should proceed.

To say that metaphors determine the ‘nature of the answers’ can be read in quite different ways, and with quite incorrect implications. Computational metaphors determine the nature of answers in cognitive science only in the sense that answers are to be framed in terms of information-processing and in principle translatable into algorithms. The *content* of those algorithms, however, is wholly underdetermined by metaphor. We can still accept a charitable reading of the idea: the fundamental generative metaphor of cognition-as-computation is what grants unity to the discourse and to a set of research problems across its constituent disciplines. It is integral to its ‘creole’ vocabulary (Galison, 1997), and implies a definitively computational character for its tools and

theories alike. In this chapter I expand on this idea of reciprocal influence between tools and theories (Gigerenzer, 1996), while arguing that both metaphors and tools can always be developed in a heterogeneous multiplicity of ways. Metaphors have enormous implications for scientific theories, but they are not the sole or even the predominant influence in the processes by which they are constructed. Rather, theories are formed at the intersections of generative metaphors with experimental setups and simulations, as supported by a whole coterie of propositions: some intended ‘literally,’ some ‘figuratively;’ some testable, some axiomatic, some merely suggestive; all often at odds with one another.

To some extent – an extent which I can only partially explore here – what I am describing is characteristic of theorization and knowledge in general. As David Leary contends, “All knowledge is ultimately rooted in metaphorical (or analogical) modes of perception and thought. Thus, metaphor necessarily plays a fundamental role in psychology, as in any other domain” (Leary, 1990, p. 2). He too is aware that he is ‘far from the first’ to propose that language and thought are fundamentally metaphorical, dating the idea back to Aristotle. Metaphor, in this tradition, is really a shorthand for any redescription of one semantic domain in terms of another.³ To use a well-entrenched conceptual metaphor, it is a matter of ‘seeing as.’ Or we may conceive of it as applying to something a name – and corresponding description – which belongs by convention to another thing. The somewhat circular nature of these attempts indicates the difficulty in constructing a ‘literal’ definition of metaphor. There are any number of other metaphors for metaphor, and a whole body of relevant work on the structure of metaphor, both in general (Arbib & Hesse, 1986; Black, 1963; Ricœur, 1993) and specifically in relation to cognition (Arbib, 1972; Gentner, Holyoak, & Kokinov, 2001;

³ This means the term is typically used to encompass other tropes such as simile and analogy, though others following Douglas Hofstadter prefer to use ‘analogy’ in the general case.

Hofstadter, 1979, 1996, G. Lakoff & Johnson, 1999, 2003; Rosch & Lloyd, 1978). All of these studies, in different ways, lend credence to the thesis that our capacity for metaphorization, broadly conceived, lies at the root of our cognitive capacities. Within language we can fruitfully understand phenomena in terms of things we know or believe about others; language and symbolic mediation in themselves can be understood as metaphorical in a more fundamental sense, inasmuch as words stand in for objects, and mathematical symbols can be manipulated in place of them.⁴ There are many layers of mediation at work in cognition. This chapter explores how metaphor runs through them and connects them to one another.

When we view metaphor as pervasive within and essential to all human thought, it of course follows that the discourse of cognitive science, like any other field of endeavour, is highly metaphorical. But this is a rather empty proposition without further specification. What kinds of metaphors are used in cognitive science? What do they *do* for cognitive scientists? How do they link together various actors and help them to cohere? What kinds of resistances and transformations do they generate as they are disseminated out into the world? These questions underpin all that follows in this work. The character of metaphors in cognitive science is quite clearly computational, and yet we shall see that this is a far from homogeneous domain. Here, though, I begin by reviewing some other relevant work on metaphor, and then some antecedent technological metaphors of mind: clocks, railways, telephones, and telegraphs. McLuhan contended that “just as a metaphor transforms and transmits experience, so do the media” (McLuhan, 1994, p. 59). Some of the very

⁴ Nietzsche is a notable advocate of this view, in a relevant passage which I cite below. McLuhan makes a similar point by playing off the common etymological roots of ‘metaphor’ and ‘translation,’ linking both concepts to mediation more generally: “All media are active metaphors in their power to translate experience into new forms. The spoken word was the first technology by which man was able to let go of his environment in order to grasp it in a new way. Words are a kind of information retrieval that can range over the total environment and experience at high speed. Words are complex systems of metaphors and symbols that translate experience into our uttered or outered senses. They are a technology of explicitness. By means of translation of immediate sense experience into vocal symbols the entire world can be evoked and retrieved at any instant” (McLuhan, 1994, p. 57).

same technological extensions of our minds which fascinated him have also proven amongst the most productive sources of metaphors for characterizing our ‘internal’ cognitive and experiential selves.

It is axiomatic to my analysis that the literal and the figurative are not two neatly separable kinds of language use. The dream of a perfectly literal language is about as realistic as the recurrent dreams of a ‘universal character’ – most notably expounded by Leibniz (Cohen, 1954) and John Wilkins, then renewed in a modern guise with logical positivism. In my view these remain two domains worth distinguishing analytically, but in practice the boundary between them is exceedingly porous and mobile. Any actually existing text is a hybrid of the two. I am more sympathetic with Nietzsche, who saw literal truth as produced from many layers of metaphor:

What is a word? The image of a nerve stimulus in sounds... The ‘thing in itself’ (for that is what pure truth, without consequences, would be) is quite incomprehensible to the creators of language and not at all worth aiming for. One designates only the relations of things to man, and to express them one calls on the boldest metaphors. A nerve stimulus, first transposed into an image—first metaphor. The image, in turn, imitated by a sound—second metaphor. And each time there is a complete overleaping of one sphere, right into the middle of an entirely new and different one (Nietzsche, 1873).

As the famous phrase from this essay has it, ‘truth’ is thus refigured as “a mobile army of metaphors, metonyms, and anthropomorphisms,” an assemblage of tropes which have become routinized to the point that we have forgotten their metaphorical character. Prefiguring later developments in the philosophy of language and sociology of knowledge, this makes truth into a “sum of human relations” (*ibid*). But while the front lines are always shifting between the figurative and the literal, it is unproductive to elide the distinction altogether.

In my view this is the wrong way to read both Nietzsche and later, like-minded scholars in science and technology studies.⁵ As Nietzsche goes on to point out, further layering on the metaphors himself, while the construction of concepts was perhaps once a merely linguistic affair, now science has taken pride of place in this endeavour. They continually labour on this ‘bulwark’ of concepts despite the fact that they can never access the ‘things in themselves’ to serve as its foundation; the scientist seeks shelter beneath the bulwark of their constructed truth, “and he requires shelter, for there are frightful powers which continuously break in upon him, powers which oppose scientific truth with completely different kinds of ‘truths’ which bear on their shields the most varied sorts of emblems” (*ibid.*). This is a quite polemical figuration which I cannot fully endorse, but has significant parallels in STS, particularly the earlier work of Latour.⁶ In my view it does suggest the best way to understand this underdetermination and undecidability. Philosophically speaking, there is no privileged access to reality. We can construct multiple systems of truth consistent with the available data, and we cannot decide *a priori* which sorts of claims are to be interpreted literally and which figuratively. This plays out most evidently of course in truths contested between science and religion. Accepting underdetermination does not imply relativism, however. Practically and sociologically speaking, what we have to do is determine and decide in the absence of any absolute standard of truth or universal decision procedure. What the ubiquitous

⁵ Haraway for instance makes the point in an interview that her fascination with the ‘cyborg’ concept is tied with the porosity of the literal/figurative divide: cyborgs are “places where the ambiguity between the literal and the figurative is always working. You are never sure whether to take something literally or figuratively. It is always both/and. It is this undecidability between the literal and the figurative that interests me about technoscience. It seems like a good place to inhabit. Moreover, the cyborg involves a physicality that is undeniable and deeply historically specific” (Haraway, 2004, p. 323). I further explore the history of cyborgs, cybernetics, and scholarly engagements with both in the following chapter – here, though, I am more concerned with this ambiguous, undecidable boundary in itself.

⁶ By framing this as a struggle exclusively between science and ‘completely different kinds of truths,’ Nietzsche’s poetic phrasing does miss the point that many of the most significant struggles occur between different propositions claiming the mantle of the real scientific truth – only *after* the resolution of such a controversy, as Latour observes, can one set of propositions be deemed non-scientific and not in accord with Nature (Latour, 1988).

polemical metaphors for such situations do rightly illuminate is that the social processes by which this occurs can often be quite antagonistic.

It is not the case, then, that the discourse of the scientist is wholly literal and unlike the ‘merely figurative’ discourse of the poet or the theologian. All discourses blend together the literal and the figurative. Yet were it not for the existence of processes by which some propositions are collectively deemed figurative, others literal, and still others as tropes conveying or ‘simplifying’ literal truths, then in its labours and social functions science would indeed be indistinguishable from poetry. This is really just another way of viewing the public processes of interpretive closure whereby ‘potentially limitless’ debates over scientific knowledge are settled in practice, however provisionally and incompletely (Collins, 1983). A proposition becoming literal as opposed to figurative, as Nietzsche framed it, is indeed partly just a matter of convention and habit—as when we cease to ‘hear’ the metaphor in referring to the end of a bed as its ‘foot,’ or the analogy between a grouse’s foot and a family ‘tree’ diagram in the old term *ped de gris* is elided by the modern ‘literal’ concept of ‘pedigree’ (G. Lakoff & Johnson, 2003). In the context of science, however, it is something more than habit and convention which carries us from metaphor to literal description. In some contexts the boundary between the two kinds of proposition is relatively well-established, but in others, particularly biology and psychology, substantial controversy endures over the right metaphors and just how far they can be taken as literal models. The mechanisms by which these controversies are settled, as I have suggested, are social, but my understanding of social is a broader one more in line with both actor-network theory and the views of human actors I study.

Such controversies are settled not by some pure, internal logic of scientific discovery—though formal logic does play a role in cognitive science to perhaps a greater degree than most disciplines, because of its close affiliation with computer science. Nor are they settled by purely

‘external’ forces such as scientists’ jockeying for power and status, or military and business interests—though these, too, play an important role in determining the paths this research follows. Rather, I agree with the researchers themselves that the most important mechanism in deciding the outcome of controversies in cognitive science is the forging of ties with non-human actors. It is a matter of what actual mechanisms one may construct, what kinds of things they can do, and how one can link them to a broader theory. I wish to complicate this story though by showing that metaphor is how these links are forged. It is by no means novel to apply technological metaphors to the mind. What is distinctive about twentieth-century cognitive science is how it was able to close the loop and follow a recurrent, recursive path from tools to theories, and back again. This phrasing of ‘tools to theories’ was coined by philosopher and cognitive scientist Gerd Gigerenzer, who argues that

(a) Scientific tools (both methods and instruments) suggest new theoretical metaphors and theoretical concepts once they are entrenched in scientific practice [and] (b) Familiarity with the tools within a scientific community also lays the foundation for the general acceptance of the theoretical concepts and metaphors inspired by the tools (Gigerenzer, 1996, p. 3).

His focus is on the mathematical tools of probability theory as they went into shaping conceptions of ‘rationality,’ but he notes that “the computer, serial and parallel, would be another case study for the tools-to-theories heuristic. . . . Physical tools, once familiar and considered indispensable, also may become the stuff of theories” (*ibid.*, 16-17). This is the case study which I undertake in the next section: how working with computers shaped conceptions of mind, and how those in turn informed the design of new programs and architectures.

Metaphors play a crucial generative and mediating function in cognitive science. They inform new ideas about how to design models and they serve as rhetorical linkages between artificial models and natural mental processes. They are not, however, a source of deterministic effects, and they by

no means fully specify the nature of scientific theories. While in practice the lines are blurry, it is crucial to maintain conceptual distinctions between more figurative and more literal claims, and between propositions of natural language and those of constructed mathematical and computational languages. As Karl Pribram puts it, analogy has been “fruitful in neuropsychology from the beginning,” sometimes explicit, “as when the brain is compared to a telephone switchboard or to the central processing unit of a computer,” but often implicit as well (Pribram, 1990, p. 81). He distinguishes though between metaphor in the general, generative sense I have been discussing, ‘analogy,’ as the process of reasoning through different facets of a metaphorical comparison in search of more literal connections, and a true *model*, “a precise coupling of an organization of data to another mode of organization such as a mathematical formulation” (*ibid.*, p.97)—or, in what amounts to much the same thing, a computer program.

The process of working from metaphor to analogy to model characterizes the history of cognitive science quite well, with the proviso that it is not in the least a continuous story of progress. Often when a particular problem or task proved intractable for a given approach, it required researchers to begin anew with a sometimes slightly, sometimes radically revised metaphor. And so while informed by the sociology of knowledge, my view of scientific theory is in line with Arbib and Hesse’s, striking a middle path between the excesses of realism and idealism: “scientific theory provides constructed models of scientific reality that are distinguished from other types of social and poetic construction by being constrained by feedback loops involving experimentation in the natural world” (p.159). In the case of cognitive science, however, these feedback loops have as much to do with experimentation in the artificial worlds of computer systems as in the (also rather elaborately

constructed) ‘natural’ world of the psychological laboratory. The question of which model is ‘correct’ is thereby supplanted to some extent by that of what you want your model to do.⁷

To focus on models and technologies in this fashion is not a shift away from broader social concerns, but a necessary step in understanding how these become inscribed *within* the tools of scientific labour. The development of cognitive science which I analyze is a key element within what Hunter Heyck has aptly characterized as the ‘Organizational revolution,’ a movement which took on mainstream status from 1955-1970, and was from the start intimately connected with the development of control technologies.⁸ This movement was tied to developments in “business, politics, and society more generally,” a turn toward formal modelling of systems with an eye to simulation (Heyck, 2014, p. 3). Whether it was of corporations, military installations, whole economies, or individual minds, such simulation was seen as embodying certain general principles and techniques.

It took place increasingly, almost exclusively by the end of this period, on digital computers. Analog control technologies did also play an important role, however, as I discuss in the next chapter. There was nothing deterministic about the effects of the computer or the metaphors it

⁷ As one contemporary brain modeling researcher frames it in a post on the pseudonymous blog *Nucleus Ambiguus*, “Whether they admit it or not, all scientists use conceptual models to organize their thinking about observations and experiments, just as all people organize their perceptions around more or less explicit understanding of how the universe works. Researchers who do mathematical/computational models (as I have) are required by the technique to specify all the moving parts in their models and to make explicit exactly how those parts interact. In that way, quantitative modeling disallows certain kinds of hand-waving (“the super-ego keeps the id in check”) while introducing others (“This symbol I wrote down means ‘supply’”)... So what makes a good model? Trick question. There is no such thing as a good model, at least not outside of the context of a particular scientific question. It all depends” (Anonymous, 2013). Issues related to brain modeling specifically are explored in greater detail within the following chapters.

⁸ Heyck specifies a useful definition of this concept (also employed by David Edge above in much the same sense): a control technology is “a device or formalized procedure that is used to coordinate the operations of multiple components so that they function as a single unit.” As he goes on to note, “Control technologies have existed for a long, long time, but the development of large-scale electric power generation and communication systems in the late nineteenth and early twentieth centuries enabled a great leap in the power and scope of such technologies, allowing the “real time” coordination of vast, integrated systems of production, distribution, and communication” (Heyck, 2014, p. 5).

generated, but as the most ‘universal’ form of control technology there was a certain inevitability to its ascent. As Heyck puts it,

Such control technologies, especially but not exclusively the electrical ones, provided a set of tangible models of the intangible structures of the world for the many who were looking for just such models. At the same time, many of these control technologies were useful not only as heuristic models of nature but also as concrete instruments for investigating, representing, or controlling it. Thus, we see some scientists using these technologies ‘at the bench,’ to investigate or represent nature, some attempting to improve these technologies through research into their fundamental properties, some using these technologies as heuristic models to guide their research, and still others not using these technologies in their research but still being influenced in the selection of problems, concepts, or methods by the programmatic goals of the new patrons of social science (Heyck, 2014, p. 5).

In a drive led by multidisciplinary scholars and ‘scholar-patrons’ like Herbert Simon, George Miller, Warren Weaver, and J.C.R. Licklider, this set of principles for understanding all levels of society through modelling rose to prominence in many disparate fields. As Heyck notes, a wide range of studies have analysed different aspects of this broader phenomenon (Haraway, 1990; Galison, 1994; Edwards, 1997; A. C. Hughes & Hughes, 2000; Mirowski, 2001; Pickering, 2011). The links are not a matter of ideas being vaguely ‘in the air,’ he argues, but of a “network of tangible connections, shared experiences with new machines and new institutions, and links to common patrons that connected a number of extremely influential people across a range of fields” (Heyck, 2014, p. 7).

Alongside the rise of new tools, this movement brought a shift in its underlying system of metaphors. The ‘organizational revolution’ of the mid-twentieth century was distinct from the earlier rise of Taylorism and like-minded movements in the social sciences which sought to maximize the efficiency of human labour. Interest had turned from exertion, fatigue, and material production to cognition, information, and symbolic production. As Heyck goes on to argue, alluding to Anson Rabinbach’s study of the earlier period, “one could describe the trajectory of twentieth-century biosocial thought quite well under the heading of a shift from the ‘human motor’ to the ‘human information processor’ as its root metaphor” (Heyck, 2014, p. 8; cf. Rabinbach, 1992). This turn is

thus intertwined with not only the rise of new tools for information-processing, but a new political and economic world. The organizational revolution in the West responded to a world in which Cold War gamesmanship had supplanted large-scale warfare, and businesses based in services, data, and media were supplanting those founded on assembly-line production. While on the one hand the computer is a generative metaphor and a means of modelling abstract phenomena like the mind, its rise to ubiquity was just as much about its real efficacy as a tool in carving out economic and geopolitical advantage for what was then known as the 'First World.' There is no one path from tools to theories, no simple line proceeding in one direction, but a multiplicity of looping paths. The courses of these paths are set not only by technologies but by metaphors and by the distinctive social problems of each era. The development of cognitive science is thus both influenced by and a crucial influence upon the shift from an economy founded upon exertion to one of attention and information. Technologies of information have become indispensable simultaneously as tools across our society and as metaphorical resources for 'objective self-fashioning' (Dumit, 2003) for understanding our selves and the nature of human subjectivity.⁹

The advent of the digital computer was hence a transformative, multifaceted event for the human sciences and for global culture more broadly. There have been many excellent studies of this to date, but as Heyck argues, "more still could be done to explore computing as metaphor, model and practice in recent science" (Heyck, 2014, p. 31). In this dissertation I undertake a further examination of a particular subdomain within this broader shift of biosocial thought: from the

⁹ This is a fundamental link between science and the problems of production and governance, noted of course by Foucault in his own way and more recently by Eve Chiapello and Luc Boltanski: "This is how the forms of capitalist production accede to representation in each epoch, by mobilizing concepts and tools that were initially developed largely autonomously in the theoretical sphere or in the domain of basic scientific research. This is the case with neurology and computer science today" (Boltanski & Chiapello, 2007, p. 104).

human motor to the human information-processor, as both metaphor and tangible model, both scientific labour and popular culture. The remainder of this chapter aims to offer some context in terms of the history of mental metaphors, both technological and otherwise. It is indeed “mind boggling to review past and present theories of consciousness and cognition with an eye peeled for metaphor. Even on the shallowest inspection, it is apparent that there have been *nothing but* metaphors in the history of these two topics” (Bruner & Feldman, 1990, p. 230).¹⁰ Yet at the same time, “many scientists feel uncomfortable with the explicit use of analogy in their work... They want to understand the results of their experiments solely in terms of those results” (Pribram, 1990, p. 79). This simply isn’t feasible though when one is attempting to relate several levels of inquiry, as in cognitive science, tying observable processes of brain physiology or computer models to the complex set of metaphorical and philosophical constructs known as ‘mind.’ In this “self-reflective process by which, metaphorically speaking, brains come to understand themselves, analogical reasoning is inevitable” (Pribram, 1990, p. 79). In a sense there is nothing unexpected about the computer analogy for mind. We have long turned to the most complex and lively technological artifacts of our own creation in search of metaphors for understanding ourselves.

Amongst the earliest modern proponents of such an analogy was Thomas Hobbes, who was quite fond of clockwork metaphors for the mind. In their heyday, “clock analogies were as popular and important for psychological theorizing as computer analogies are today” (McReynolds, 1980, p. 97). Hobbes aimed to go beyond such analogies, however, arguing that all ‘automata’ truly have an

¹⁰ Bruner and Feldman go on to enumerate some examples: “spotlight, footlight, flowing river, stream, seamless web, graph, recursive loop, pandemonium... there is no end to this parade” (Bruner & Feldman, 1990, p. 230). As they go on to note there has been a particularly longstanding division between two metaphor-systems, one which has consciousness developing from the inside out, the other from the outside in (*ibid.*, p. 233); once, this was played out as a dispute between empiricism and rationalism, today, between proponents of cognition as symbol-manipulation and those arguing for cognition as situated, embodied, and action-oriented (Agre, 1993, 1997; Suchman, 1987).

‘artificiall life’ about them (Hobbes, 1651, p. 1).¹¹ The core of the *Leviathan*, really, is its exposition of this tripartite analogy linking the ‘life’ of clockwork-type mechanisms to the structure of minds, states, and the world as a whole: just as the designer of automata imparts purpose to an assemblage of cogs and springs, so too does God order nature, the sovereign the state, and the rational consciousness our perceptions and ideas. This is a process which can go wrong in myriad ways. Nature can produce monsters, leaders can govern poorly, and consciousness is subject to illusions and errors. For Hobbes, though, there is only one proper fashion to proceed in the sphere of human reason and action, and that is on the model of geometry, mathematics, and logic – by careful argumentation constructed on the basis of axiomatic definitions. This was the art of logic in his view, ‘adding’ and ‘subtracting’ propositions to produce syllogisms and sound demonstrations in agreed-upon ways. Such was the combination of mathematics and mechanism in his conception of human nature that more recent philosophers of mind have labelled him the “grandfather of AI” (Haugeland, 1989, p. 23). He did directly equate “ratiocination” with “computation” (*ibid.*), and though evidently anachronistic the phrase ‘artificial intelligence’ would indeed fit quite well with the vocabulary of the *Leviathan*.

Hobbes was not the only one in his time to relate human thought and behaviour to clockwork mechanisms. Descartes did likewise, contending in *The Passions of the Soul* that the movements of both humans and animals were governed by “the brain, nerves and muscles, just as the movements of a watch are produced simply by the strength of the springs and the form of the wheels” (in McReynolds, 1980, p. 100). On his view, animals could be fully understood in

¹¹ Like many other philosophers Hobbes was skeptical of metaphor despite using it profusely throughout his work, insisting that the use of “Metaphors, Tropes, and other Rhetorical figures instead of the words proper” – that is to say words as carefully defined in accordance with Hobbes’ principles – was a cause of “absurd conclusions,” proceeding from the “confusion” and “unfit connexion” of words and concepts (Hobbes, 1651, p. 20).

mechanistic terms: as assemblages of muscles and nerves whose behaviours were driven by instincts and reflexes.¹² Designers of automata like Jacques de Vaucanson constructed elaborate demonstrations of organic bodies replicated in machinery, as in his famous duck which quacked, ate, flapped its wings, and defecated, along with his lesser-known flute-playing automata. Though some later commentators have denounced these as mere showmanship, they were in fact intended as serious experimental works in physiology (*ibid.*, pp.81-87). Really, they were all of the above: showpieces, entertainments, and experiments, domains which remained closely intertwined in the art of automata-building on into the nineteenth century (Schaffer, 1997).

Descartes saw behaviour and some aspects of perception as mechanistic and physicalistic, but drew a line at human consciousness and cognition. For him, the clockwork analogy did not extend to mind. Hobbes and more vociferous advocates of mechanism like Julien Offray de la Mettrie saw no problem with this. As the latter argued, “since all the faculties of the soul depend so much upon the proper organization of the brain, and of the whole body, that they appear evidently to be nothing but this organization itself; we may well call it an enlightened machine” (Offray, 1748, p. 47). Perhaps we have been granted “some more wheels, some more springs, than are found in other animals” (*ibid.*), but we are machines all the same. Against this increasingly insistent deployment of clockwork analogies, however, others continued to reject them as Descartes had. Leibniz was still more skeptical of their applicability even to animal physiology, saying that the ‘unity of a clock’ is something “quite other than that of an animal; for an animal may be a substance possessing a genuine unity, like what is called ego [*moi*] in us; while the clock is nothing but an

¹² He even offered a speculative explanation of how this occurred physiologically, based on his viewing of hydraulic automata in the royal gardens, suggesting that like the pipes within those mechanisms nerves perhaps channeled a sort of fluid (Descartes, 2003).

aggregate [*assemblage*]” (in McReynolds, 1980, p. 102). Technological metaphors thus set the terms of debates about the nature of the soul, though they did not by any means determine their outcome.

Such debates have in principle been set aside by current science, but those regarding the nature of mind present the same fundamental problem in secularized form. Whether it be labelled soul, mind, or consciousness, is there something which *transcends* and directs the physiological mechanisms of our body? Or is our experience of subjective agency and free will something that at best emerges *from* underlying mechanisms? Is mind perhaps even an illusory epiphenomenon, merely echoing rather than directing the real causal processes at work? Technological metaphors are now more closely allied with the latter view, but this was less clear in the case of clockwork metaphors. They lent themselves just as readily to Cartesianism, which sought to draw a firm distinction between those aspects of the physical world which could be explained in mechanistic terms, and the interior world of human cognition, which could not be. Even otherwise staunch advocates of the metaphor had to acknowledge that humankind was rather unlike a clock in our ability to ‘wind our own spring’ (Offray, 1748). The ordering and driving force of the clock is external. This differs significantly from later electrical technologies. Perhaps we simply reach for the most complex technology available in search of analogies to describe the workings of our minds, but the nature of the technologies alters the character of the metaphors. In clockwork mechanisms there is a sharp distinction between design and operation, between operative mechanisms and motive force; these diminish to some degree in telegraphy and disappear entirely in computing, as electricity simultaneously powers and controls, stores programs and carries out operations.

These features make electrical technologies immediately appealing as metaphors for cognition. Rendering the metaphor still more compelling, science began to demonstrate that it was this same electricity which coursed through our nervous system and directed the functions of our

body. This leads to the premise of cognitive science: that the metaphor points to a true similarity of underlying structure, with specifically computational technologies. First, though, came the ascent of telecommunications technologies as our primary source of metaphors for thinking about thinking. Already in 1851 Emil du Bois-Reymond was contending that “the similarity between the two apparatus, the nervous system and the electric telegraph, has a much deeper foundation. It is more than similarity; it is a kinship between the two, an agreement not merely of the effects, but also perhaps of the causes” (in Otis, 2002, p. 106). In parallel with the development of the electric telegraph – which was exceedingly rapid, from Morse’s first line from Washington to Baltimore in 1844, to the 1861 completion of telegraph lines traversing the American continent (McLuhan, 1994, p. 250) – came a research programme by du Bois-Reymond, Volta, and others to confirm that it was indeed electricity which carried impulses through organic nerves (Lenoir, 1986; Otis, 2001). The telegraph was simultaneously a driver of change in commerce, publishing, and politics; a crucial control technology for other systems, above all the railroad; the basis for new organic metaphors describing society and its transformations in this new era; all in addition to its role as generative metaphor and experimental model for a new understanding of the organism. It inspired magnates and preachers alike to wax poetic about its potential. In sum, it was the first great exemplar for the co-constitution of science, technology, and communications media, in a globe-spanning, symbolically-charged, ‘sublime’ system.¹³

The extension of human capacities through telegraphy was itself transformative, first of all for the self-understanding of those engaged most directly with the system: “Telegraph operators perceived their keys and wires as extensions of their hands and minds. At times, their hands and

¹³ This notion of ‘technological sublime’ comes from Leo Marx, who also makes a compelling case that the conjunction of telegraphy and railways in fact occasioned the rise of the very concept of ‘technology,’ to name this new configuration of scientific knowledge and what had previously been known as the ‘mechanical arts’ (Marx, 1997).

minds seemed like extensions of their wires and keys” (Otis, 2001, p. 231). Hemingway recounted being ‘fascinated by the lingo of the cable’ during his time as a reporter, and recognized the pared-down, literally ‘telegraphic’ style of writing for cable transmission as an influence on his own literary style (Carey, 2008, p. 163).¹⁴ From the beginning, alongside these popular consequences of the new medium, and the real scientific labour going into designing and maintaining its infrastructure, telegraphy was being explored as an analogy for cognition. As explored in detail by Timothy Lenoir, among the most compelling advocates for this analogy was Hermann Helmholtz. “From as early as 1850 he drew analogies between the electrical telegraph and the process of perception. The telegraph began to serve as a generalized model for representing the processes of sensation and perception” (Lenoir, 1994, p. 186), and so he came eventually to view the entire nervous system as a “media apparatus,” the eye a ‘photometer’ and the ear a ‘tuning-fork interrupter with attached resonators’ (*ibid.*, p.185). The inputs of these elements were then translated into nervous impulses through a cognitive analogue of Morse code. Through experimentation, Helmholtz sought to decipher this coding. Much as in later cybernetics and cognitive science, this involved both the systematic study of human perceptual capacities (discriminating colours or tones, for instance) and the construction of artificial models.

Electricity gives a new force to technological metaphors and models of mind. It carries connotations of the ineffable, the immaterial, of a bolt from the heavens. Running through organic

¹⁴ Much more could be said about the social consequences of telecommunications, and moreover the ways these have been interpreted by various figures. McLuhan offers many compelling examples, including the ominous later observation by Albert Speer that “former dictatorships needed collaborators of high quality even in the lower levels of leadership, men who could think and act independently. In the era of modern technique an authoritarian system can do without this. The means of communication alone permit it to mechanize the work of subordinate leadership. As a consequence a new type develops: the uncritical recipient of orders” (from Hjalmar Schacht, *Account Settled*, p.240, cited in McLuhan, 1994).

nerves, it accomplishes a multiplicity of functions, just as it can be harnessed for an ever-increasing number of artificial applications. In Helmholtz's elaboration of the analogy:

“Nerves have been often and not unsuitably compared to telegraph wires. Such a wire conducts one kind of electric current and no other; it may be stronger, it may be weaker, it may move in either direction; it has no other qualitative differences. Nevertheless, according to the different kinds of apparatus with which we provide its terminations, we can send telegraphic dispatches, ring bells, explode mines, decompose water, move magnets, magnetise iron, develop light, and so on. So with the nerves. The condition of excitement which can be produced in them, and is conducted by them, is, so far as it can be recognised in isolated fibres of a nerve, everywhere the same, but when it is brought to various parts of the brain, or the body, it produces motion, secretions of glands, increase and decrease of the quantity of blood, of redness and of warmth of individual organs, and also sensations of light, of hearing, and so forth” (in Lenoir, 1994, p. 207).

Helmholtz was far from the only advocate of this analogy in his day. In his 1855 *The Senses and the Intellect*, Scottish philosopher Alexander Bain – not to be confused with the later pioneer of telegraphy – mixed traditional metaphors of rivers and trees in the mind with those likening it to ‘the course of a railway train,’ or a telegraph in which the network of wires ‘might be formed to represent exactly what takes place in the brain’ (in Coleman & Fraser, 2011, p. 110). Even in metaphor, the railway and the telegraph were closely allied, and the former figured nicely as the ‘musculature’ to the latter’s ‘nervous system.’ In many ways then telegraphy seemed a far superior metaphor for mind as compared with clockwork.

There was of course opposition to this metaphor as well, notably from English writer George Lewes, and from pioneering Spanish neuroscientist Santiago Ramón y Cajal. Particularly for the latter, the telegraphic analogy did a “grave injustice to the nervous system by denying its plasticity, its most essential attribute” (Otis, 2001, p. 80). Despite its many potential functions, the structure of the telegraph network remained fundamentally static, and externally imposed. No matter how many messages were transmitted from one station to another, the connections would never become ‘stronger’ – quite the contrary, they might hit their capacity. Building more lines required the

intervention of human designers, with substantial labours of stringing cables, implementing amplification and decoding of signals, training operators, and so on. In these respects the telegraph system seemed to have rather little in common with the brain or the nervous system.

Other technological metaphors soon came along. Jean-Marie Guyau, in his 1880 “Memory and Phonograph,” concurs that “the greatest mystery of brain mechanics has to do with dynamics—not with statics,” and so perhaps “the most refined instrument (both receiver and motor in one) with which the human brain may be compared is perhaps Edison’s recently invented phonograph” (in Kittler, 1999, p. 30). The phonograph transfers the “vibrations of one’s voice... to a point that engraves lines onto a metal plate that correspond to the uttered sounds.” These are “more or less deep,” depending on the sounds, and he argues that “it is quite probable that in analogous ways, invisible lines are incessantly carved into the brain cells, which provide a channel for nerve streams” (*ibid.*, pp.30-31). As new thoughts and perceptions flow through the nerves, he suggests they somehow activate memories of similar past thoughts, just as the phonograph reproduces sound from tracings on a disk.¹⁵

By the early decades of the twentieth century, the techniques of sound recording and electronic signal transmission had been fused in the development of telephony. As the telephone entered wide usage, it began to furnish the preferred metaphors for the brain and nervous system. The switchboard in particular was a common referent. Originally these were literal plugboards staffed by human operators, who would see a terminal light up when an outbound call was made, plug in a cable, ask the person on the other end where they’d like their call directed, and connect the other end of the cable to that party’s jack. Later this function was automated, with electromechanical

¹⁵ After reproducing this little-known essay in full, Kittler goes on to adduce many further examples of the phonograph as both metaphor for the mind, and tool for scientific research. Myriad others are covered in his *Discourse Networks 1800/1900* (Kittler, 1990, 1999).




Figure 1 - Visual analogy between spine and telephone system

stepping switches translating the rotations of subscribers' dials into a programme for routing the call directly. Henri Bergson in *Matter and Memory* likened the brain to a 'central telephonic exchange,' emphasizing that he saw it as a mere intermediary, routing signals in various ways, sometimes allowing communication, delaying it or blocking it off, but not really transforming what it transmitted (Bergson, 1988). The telephone metaphor was also one of several technological metaphors adopted

by Pavlov in characterizing the nervous system, with inborn reflexes likened to direct, permanent lines and conditioned reflexes to the "flexible temporary connections between telephone users through a switchboard." Just as the telephone exchange "solved the problem of communication for a large number of users, the mechanism of conditional reflexes, in Pavlov's view, solved the problem of the organism's reaction to diverse stimuli" (Gerovitch, 2002, p. 344). J.S. Gray affirms much the same, noting that he and several of his contemporaries view the nervous system as a "switchboard mechanism' which merely distributes neural flux to the various motor points" (Gray, 1935, p. 111).

Others concocted fanciful diagrams likening the spinal column to a telephonic exchange (Fig. 1; Keith, 1920, p. 258).

For many scientists, then, the telephone switchboard metaphor served to emphasize the role of reflexes in behaviour and cognition, characterizing nerves as neutral intermediaries binding together stimulus and response. Particularly in its automatic variant, the telephone exchange was a powerful trope for undercutting mystical and dualistic notions. As a complex electromechanical system becomes capable of more and more sophisticated functions without requiring human intervention, it suggests that we need not invoke a transcendental soul or mind in explaining behaviour, for much of that can be likewise reduced to the paths taken by electrical signals. This was of course more plausible in the domain of specific motor responses than in more complex cognitive tasks. Karl Lashley invoked the telephone metaphor to obliquely criticize the Pavlovian view of the organism. On this account,

The performance of a habit, whether of speech or of manipulative movement, is determined by the existence of definite connections between a limited number of nerve cells, which are always functional in that habit. The model for the theory is a telephone system. Just as two instruments can be connected only by certain wires, so the sense organs and muscles concerned in any act are connected by nerve fibers specialized for that act (Lashley, 1930, p. 4).

Lashley contended, however, that there was little evidence for the ‘definite specialized synapses’ required by this theory, and that it was unwise to suppose that “the mechanism of cerebral function is essentially the same as that of the spinal reflexes” (*ibid.*). His claims were based largely on the plasticity of brain function: neural processes in the brain are not highly localized or specific, and following lesioning experiments in animal models, behaviours are often quite resilient. Rats with various seemingly essential connections severed could often find their way through mazes with only moderate degradation in performance. The telephone metaphor was thus again not uncontested, nor was there consensus about its implications.

Freud, too, employed the telephone as one of several technological metaphors in his writing, taking it in a rather different direction. Summarizing the analyst's technique, he stated that

To put it in a formula: he must turn his own unconscious like a receptive organ towards the transmitting unconscious of the patient. He must adjust himself to the patient as a telephone receiver is adjusted to the transmitting microphone. Just as the receiver converts back into sound waves the electric oscillations in the telephone line which were set up by sound waves, so the doctor's unconscious is able, from the derivatives of the unconscious which are communicated to him, to reconstruct that unconscious, which has determined the patient's free associations (Freud, 1999, p. 2470)

Telecommunications was thus taken as an equally apt metaphor for psychoanalysis as for physiologically-oriented psychology. In this case though rather than binding together stimuli and responses like a switchboard, the simile implied that just as a telephone receiver could convert words into electrical signals, the analyst could transduce psychic energies from the statements of the patient. For Freud it was signal processing rather than switching which served as the key feature of this technology in service of his metaphorical redescription. He also rejected the claims of his contemporaries that the increasing prevalence of nervous complaints could be attributed to "the immense extension of communications which has been brought about by the network of telegraphs and telephones that encircle the world," and which had "completely altered the conditions of trade and commerce" ("W. Erb" in Freud, 1999, p. 1949). Against this account, which he contended was "insufficient to explain the details in the picture of nervous disturbances," he offered his own favoured hypothesis, that "the injurious influence of civilization reduces itself in the main to the harmful suppression of the sexual life of civilized peoples (or classes) through the 'civilized' sexual morality prevalent in them" (Freud, 1999, p. 1950).

Ultimately telecommunications came to mean much more than simply metaphor for psychology. The field's most important contributions to neuropsychology "came in the form of *techniques* for measuring the flow of signals" (Pribram, 1990, p. 81 emphasis added). The most crucial

figure in this regard was Claude Shannon, who developed the first mathematical treatment of information following on from his work during wartime on cryptography and control technologies, then his later work with Bell Labs. First published in the 'Bell System Technical Journal,' his 'Mathematical Theory of Communication' sought to separate the 'engineering' dimensions of mediated communications from semantics (Shannon, 1948). This was a practical technological problem subjected to analysis in its most general form. To Shannon the problem was essentially the same whether channels were conveying Latin characters, modulations of sound waves, or the simplest unit of information, the 'one' or 'zero' of binary digits or 'bits'—the working vocabulary of both telegraphy and digital computing. To communicate successfully was to correctly and exclusively convey *one* message selected from a set of possible messages (Shannon, 1948, p. 379). This was an important problem in telecommunications of course: how noisy could a line be before speakers would be unable to understand one another? How could telephone circuits be designed to reduce that inevitable noise without also muffling the voice signal? Yet as Shannon developed the theory and as it was framed in a well-known introduction by Warren Weaver, it was of extremely general import. Shannon saw a close connection between his wartime and later work, noting that the problem of separating signal from noise in telecommunications was closely analogous to that of separating 'tracking errors' in fire-control systems from the desired 'signal' of an anti-aircraft gunner leading a target (Mindell, 1995).

Along with Weaver, another champion of Shannon's theory was Norbert Wiener, who had independently begun developing similar ideas and mathematical techniques after his own work on fire-control systems. According to his introduction of the concept, "just as the amount of information in a system is a measure of its degree of organization, so the entropy of a system is a measure of its degree of disorganization; and the one is simply the negative of the other" (N.

Wiener, 1961, p. 11). He saw the theory as bridging fundamental questions about the nature of the physical universe with those concerning organic life and its islands of ‘negative entropy’ standing against the overall tendencies of the universe. All this was in addition to its immediate relevance for engineering purposes. Interestingly, Shannon and Weaver placed ‘opposite signs on their formulations of entropy.’ The former saw information *as* entropy, not as its inverse, the way Wiener did (Gleick, 2011). Shannon called this a ‘mathematical pun,’ stating that they get the same answers despite their differing usages of the term: “I consider how much information is *produced* when a choice is made from a set—the larger the set the more information.” Wiener instead saw this larger uncertainty as meaning “less knowledge of the situation and hence less information” (letter of Shannon to Weaver, 13 Oct. 1948, cited in Gleick, 2011). For Shannon the most highly entropic message contained the most information. Conceived as a series of selections from universes of possible messages, a collection of random characters or audio tones is far denser with information, inasmuch as none of the elements within the series gives any clue as to the next. This makes a great deal of sense in terms of the meaning-agnostic, engineering approach which Shannon adopted, and his mathematical formulation became the canonical version of information theory.

In what has become an overriding concern in modern computing, such random signals are impossible to ‘compress,’ unlike strings of text or recorded speech in English or another natural language. In those cases each element of the signal does offer some indications as to which other possibilities are more likely to follow. This knowledge can be used to design information systems to maximize the utility of each ‘bit,’ as demonstrated early on in the usage of shorter codes for more frequently-occurring letters in the Morse code. It can also be used to ‘smooth’ out elements of noise automatically, by filtering out elements of an input which are unlikely to be part of the desired signal: be it literal noise, marked by sudden changes of pitch in an audio signal, or sudden ‘twitchy’

movements in the input of a fire-control system, which information theory came to understand as figurative noise. Wiener's formulation does thereby retain more intuitive appeal than Shannon's. Though random, entropic signals are far harder to compress from an engineering standpoint, from the inescapable standpoint of human beings who seek *meaning* in messages, a text or an utterance or a video strikes us as containing more information than pure noise. For that matter, our own genome likewise seems much denser with information as compared with any random collection of organic compounds, as do the coded firings of neurons compared with arbitrary fluctuations of electrical potential.

Regardless of this ambiguity, the mathematical tools and conceptual vocabulary of information theory eventually became central to cognitive science. Indeed, some early practitioners in this interdisciplinary field called it simply "information-processing psychology" (Miller, 2003). Along with Pribram, Newell, and Simon, George Miller was one of many who saw these tools as valuable resources in the development of a true science of mind, one which would no longer bracket the very concept as metaphysical or illusory. This 'mind' was a somewhat immaterial and abstract sort of entity, but one nevertheless susceptible to mathematical analysis and technological manipulation, much like the messages borne by electronic telecommunications. I examine this history in detail in the next chapter. Information theory provided at once a source of practical techniques among the various disciplines which made up cognitive science, as well as of metaphors and a common figurative vocabulary for its process of 'creolization' (Galison, 1997). At the same time it did more to 'sharpen the framing of questions' than to provide specific answers. (Pribram, 1990, p. 82). The 'control technologies' from which the theory was borne, and whose later designs it went on to inform, furnished tangible models and popular metaphors for the intangible realm of cognition. They became tools for comprehending a secularized version of the soul, conceived along

Hobbesian lines. Technological metaphor is thus a crucial part of the story here, but so too is the labour of constructing actual devices to observe how they support or resist such analogies.

Cognitive science carried this entwined theory and practice to new heights in the twentieth century, but this was equally what drove figures like de Vaucanson or Babbage. As technologies changed, so too did our metaphors and theories of mind. As Slava Gerovitch observed in his study of Pavlov and Soviet cybernetics, “man-machine metaphors travel by a spiral rather than a circle: at each stage new, more complex machines provide metaphors for more sophisticated physiological concepts, and vice versa.” The old metaphors and their associated technologies are often seen as “‘mechanistic’ and reductionist” in negative ways, even as new ones “seem liberating by overcoming some limitations of the old” (Gerovitch, 2002, p. 369). This is very much the case with computational metaphor in cognitive science. Now it has come to bear some of the same connotations of rigidity and dehumanization as the earlier clockwork-style mechanism of Hobbes, depending whom one asks. Yet for cognitive science these new electronic control technologies represented machines that seemed finally up to the task as metaphors and models for human mind and behaviour. They were driven by the same fundamental force as that which ran through our bodies, they seemed capable of just as many diverse actions and responses, thus building them in our image might be of both practical and theoretical value. We could build computer programs to assist us by performing cognitive tasks, and in so doing better understand what cognition really was.

Metaphors structure and guide the course of scientific research, as they do public discourse, but by no means do they determine either entirely. Science, rather, involves a continuous bidirectional traffic between analogical creativity and mathematical specificity, between the heights of logical purity and the messy social circumstances in which its labour always takes place. This is particularly evident in the case of cognitive science, as it traces a recurrent and recursive path from

tools to theories and back again. What makes computation so powerful, as I explore in the next chapters, is that it can function simultaneously as generative analogy, technological supplement, and tangible model for the mind. Computers are ‘things to think with’ in every sense, at once indispensable tools for research, figurative spurs to insightful connection-making, and the artifacts which can most plausibly be said to ‘think’ for themselves. In their early days, they served as the focal point for a social collective and core set of researchers who would come to define themselves as cognitive scientists. Computers became a focus both intellectually and materially, as they would often find themselves congregated around these massive machines when processing time could be diverted from more pressing, typically military uses. Often employed at those same pursuits during the day, the pioneers of cognitive science could then be found ‘hacking’ away in the wee hours of the morning on their own attempts to, as Babbage phrased it, ‘throw the powers of thought into wheel-work.’

3

A mediated history of cognitive science

“This, at least, is my reading of what motivated the pioneers of cognitive science: the notion that thought, mental activity, this faculty of mind that has knowledge as its object, is in the last analysis nothing other than a rule-governed mechanical process, a ‘blind’—might one go so far as to say ‘stupid’?—automatism. Did this amount to devaluing humanity? To elevating the machine? Or, to the contrary, did they see man as a demiurge, capable of creating an artificial brain or mind? No doubt there is some truth to each of these interpretations, more or less depending upon which individual scientists are considered and the times in which they worked.”

— Jean-Pierre Dupuy (2000, p.39)

In three closely linked chapters, this section offers a brief history of cognitive science and its constituent disciplines, developing a case for their profoundly *mediated* character. Writing the history of this relatively young interdisciplinary field is challenging, and so I employ the concepts of media and mediation as structuring devices, arguing for their specific relevance in this area. There are of course distinct histories of psychology, psychiatry, and neuroscience which overlap throughout, and which have already been studied in depth from a variety of perspectives, often considering similar questions of metaphor and mediation (Danziger, 2008; G. Lakoff & Johnson, 1999; Lenoir, 1986, 1994; Otis, 2001). The history of computer science and artificial intelligence (AI) also figures prominently in the genesis of cognitive science, and has been well chronicled at least up to the 1990s (Crevier, 1993; Ekbia, 2008; McCorduck, 2004). Many actors in these fields have also offered their own valuable historical accounts.¹ All of this valuable prior research informs the discussion which

¹ These are mostly written in an internalist mode, but nevertheless constitute valuable evidence of the field as perceived by researchers within the different periods I discuss (Newell, 1982; Dupuy, 2000; Miller, 2003). More recently, in all of

follows. I only consider these disciplines, however, in their intersection with the history of cognitive science proper, which in my view has yet to be articulated in an entirely satisfying way.

The field is so heterogeneous as to defy any easy summary, and many have observed that it would be preferable to speak of the *cognitive sciences* in the plural. From its origins, it was envisioned as essentially transdisciplinary and interdepartmental. As Jerome Bruner and George Miller wrote in a letter to their dean about the founding principles of the Harvard Center for Cognitive Studies, “The slogan, only half in jest, was that the cognitive processes are far too complex and important to be left to psychologists” (in Cohen-Cole, 2007, p. 568). Their Center was one of the key sites in the ‘cognitive revolution,’ the well-known story of behaviourism’s decline and the return of ‘mentalistic’ concepts to scientific respectability (Miller, 2003; Miller, Galanter, & Pribram, 1960).² Yet as Miller himself recalls, this redefinition of psychology represented only part of the story. The other crucial contemporaneous developments were the rise in popularity of Norbert Wiener’s cybernetics, the conceptualization of AI by Marvin Minsky and John McCarthy, the closely related use of computers for ‘cognitive simulation’ by Alan Newell and Herbert Simon, and finally the redefinition of linguistics by Noam Chomsky (Miller, 2003, p. 142).

This partial list of protagonists gives an idea of the complexity in coming to terms with cognitive science. Another indication comes from the remarkable length of the one attempt at a comprehensive history of the field, Margaret Boden’s two-volume *magnum opus*, *Mind as Machine*

the affiliated disciplines, the Internet has also given rise to a proliferation of actors’ own accounts of the history and ongoing development of their research programme. I draw on some of these toward the end of this chapter, and as I discuss at greater length in the final chapter, one of the myriad consequences of this medium for scientific discourse is that the lines of ‘appropriate’ professional debate are blurring. Controversies or polemics that would have been matters of informal ‘corridor talk’ and private letters are now quite public, available to STS scholars (among many other audiences) in real-time rather than through archival bequests decades later.

² As Miller recounts, “We were still reluctant to use such terms as ‘mentalism’ to describe what was needed [for psychology to succeed], so we talked about cognition instead” (Miller, 2003, p. 142).

(Boden, 2006).³ ‘Cognitive science,’ in its singular or plural form, stands today as one of the preeminent self-definitions for the scientific study of mind.

Neuroscience is also of course a powerful touchstone, but many neuroscientists pursuing cross-disciplinary work situate themselves within the umbrella discipline of cognitive science, whether informally, or through participation in the Cognitive Science Society, or in similarly labelled academic programs. What unifies all of




Figure 2 – Constituent disciplines of cognitive science

the disparate endeavours gathered under this umbrella, though? What value does this cross-disciplinary categorization have? Another part of the story, to be sure, has to do with the imperatives of funding. Otherwise, there might have been an endless proliferation of distinct formulations rather than one overarching field known as cognitive science. Newell and Simon say that ‘1956 could be taken as the critical year’ for the development of cognitive science (Newell & Simon, 1972, p. 878): all the key actors had begun to communicate with one another, and Miller dates the ‘moment of conception’ even more specifically to the convening of a special interest group on information theory at MIT on September 11, 1956. It was clear ‘by 1960’ that ‘something interdisciplinary was

³ Boden’s work is immensely valuable, and in many ways my analysis proceeds in a similar fashion, as a history of technological metaphor in the science of mind. But unfortunately it does not render the history of cognitive science as such in an especially satisfying way. In the first place, it is so large and unwieldy, covering such a broad scope of material that one seeking a clear understanding of the general course of research and specific disciplinary identity of the field might just as well consult the archives and prior literature themselves. (Comprehensive, lengthy histories of science need not be so unwieldy, as illustrated by the clarity of structure in, for instance, Galison, 1997). More problematic is the polemical agenda, rarely made explicit but running throughout the entire work, against Chomsky’s approach to linguistics, which has raised questions about the accuracy and validity of Boden’s account in parts (Casper, 2010; Chomsky, 2007). The other historical account which aims to capture the full scope of the ‘cognitive revolution’ is Howard Gardner’s more straightforward and accessible *The Mind’s New Science* (Gardner, 2008). I draw on this work as well, but focus on tools rather than disciplines or theories, and I endeavour to disentangle cognitive science proper from the eclipse of behaviorism and the broader cognitive revolution.

happening.’ at Harvard they called it ‘cognitive studies,’ at Carnegie-Mellon (home of Newell and Simon) it was ‘information-processing psychology,’ and at La Jolla it was ‘cognitive science.’ But “what you called it didn’t really matter until 1976, when the Alfred P. Sloan Foundation became interested” (Miller, 2003, p. 143). The diagram above is reproduced from Miller’s 1978 report to the Sloan Foundation on the state of cognitive science, reflecting their understanding of its constituent disciplines and existing lines of interdisciplinary inquiry at the time. Having just completed a successful funding program in the budding neurosciences, the Foundation considered the next step to be bridging the gap between mind and brain, and in dialogue with Miller and a committee of like-minded researchers, they settled upon the name of ‘cognitive science’ for this pursuit.

What, then, defined these lines of interdisciplinary inquiry? In a sense it was their interest in ‘mind.’ This begs the further question, however, of what factors contributed to the concept’s renewed scientific respectability. I follow Jamie Cohen-Cole’s focus on the kinds of *tool-exchange* that occurred within cognitive science rather than the particular theories, institutions, or departments involved (Cohen-Cole, 2007). These tools, with their international patterns of exchange, use, and modification, are the only unity we find across the history of this ‘pluralistic’ field. Such tools took many forms – equally including instruments for physiological measurement and for conducting surveys – but in this section I focus chiefly on two kinds of interrelated tools, namely the techniques of mathematization borne of information theory and cybernetics, and the hardware tools of computing, which rendered such mathematics tangible and semi-autonomous. I suggest that the ‘intellectual and social world’ not only of the Center for Cognitive Studies but of cognitive science in general is captured better by these economies of tool-exchange and –use amongst researchers than by any theory-driven sense of a ‘paradigm’ (Kuhn, 1957). Metaphors drawn from computing and control technologies constitute the clearest common thread tying together the heterogeneous

	<i>Approach</i>	<i>Approximate Dates</i>	<i>Characteristic technologies</i>
approaches which came to be defined as cognitive science.	<i>Cybernetics</i>	1948-1969	Weapons control, Homeostat, Perceptron
Extending this argument by way of Gigerenzer's 'tools-to-theories	<i>Symbolic CS/ AI ('GOFAI')</i>	1956-1974	Logic Theorist, General Problem Solver, LISP
heuristic,' as discussed in the	<i>Post-symbolic, parallel & hybrid</i>	1986-today	PDP, Blue Brain, DARPA Synapse

previous chapter, I argue that the field is constituted by an ongoing process of feedback and translation, first of all between machines and ideas, but also between two modes of engagement with its fundamental equation of mind-as-computation: from suggestive metaphor to working model—and back again, in an ongoing reciprocal interaction. This process is at the heart of what I call 'mediated cognition:' the profound, distinctive coupling between minds, theories, and technological media cultivated within cognitive science.

For the sake of analytical clarity, I divide this history of twentieth-century cognitive science into three roughly separable, yet overlapping and interwoven strands. The *dramatis personae* given by Miller above indicates how the roots of cybernetics and symbolic, computational cognitive science were in a sense contemporaneous. Nevertheless, these can also be seen as distinct approaches to the study of cognition which rose and fell in a rough historical sequence. For each, though an element of arbitrariness is inevitable, I have selected a date of origin, marking some significant, literally paradigmatic and programme-defining work in the field.⁴ These correspond with actors' own accounts of key moments: the release of Norbert Wiener's *Cybernetics*, the Dartmouth conferences

⁴ I use the term paradigmatic here not in its special Kuhnian sense, but follow the earlier meaning which informed Kuhn's choice of terminology, that of a concrete model taken to exemplify certain key features of a given reality. There are several paradigms in this sense within the history of cognitive science, each of which played an important role in defining the course of subsequent research: Norbert Wiener's *Cybernetics*, Ross Ashby's Homeostat, or Newell and Simon's *Human Problem Solving* and their General Problem Solver architecture. These three eras are not separate and 'incommensurable' Kuhnian paradigms, however. Though full engagement with the debates around Kuhn in STS is beyond the scope of this essay, it does seem that cognitive science forms one coherent paradigm or research programme (Kuhn, 2012; Lakatos, 1968) centred around the treatment of mind as information-processing and computation.

on artificial intelligence, and the revitalization of connectionism with David Rumelhart and James McClelland's *Parallel Distributed Processing*. The periods which follow each, and the dates I choose to mark their conclusions, are not a matter of definitive falsification for any one approach (though I discuss the case of the Perceptron, which comes closest to that archetype). Rather, each rose and fell as a highly mediated discipline in another sense: the more intuitive way that any discipline which promises scientific knowledge about the nature of our minds commands public interest and press coverage.

As epitomized by Newell and Simon's claims in their 1958 address to the Operations Research Society, that "there are now in the world machines that think, that learn, and that create" (Simon & Newell, 1958, p. 8) and that within a decade computers would be chess champions, acclaimed composers and groundbreaking mathematicians, there was often considerable bombast, drama, and hype around claims in the field. Most focused on the dramatic claims for artificial intelligence, but in fact Newell and Simon were most interested in the use of these programs for modelling human psychological and social processes. They predicted, likewise, that within a decade most theories in psychology would take the form of computer programs, or qualitative statements about computer programs (*ibid.*). Like their other predictions, this one about the radical transformation of psychology in its entirety was off the mark; with colleagues like Miller, Chomsky, Wiener and myriad others, however, they were well under way to constructing cognitive science as a transdisciplinary field in which computer programs were not mere tools, but indeed theory-constitutive. The prospect of understanding and replicating the human mind seems to command an intrinsic allure and power of *interessement* (cf. Callon, 1986). When the tools of each era failed to deliver on their creators' over-ambitious promises, however, both public interest and public funding waned.

In each period, I focus on a few characteristic technologies used as models of cognition, allowing these more than their designers or theories to structure my historical narrative. This history is best understood in terms of a shift from biologically-inspired analog models in the era of cybernetics, to the abstract and digital symbol-manipulating models favoured by the likes of Newell and Simon, to a still more pluralistic contemporary situation. Alongside symbolic approaches and newer hybrid forms of parallel processing, we are now witnessing a large-scale return to biology, simulated on dramatically more powerful digital computers. Throughout this chapter, I situate the processes of interdisciplinary tool-exchange and the tools-to-theories heuristic as central to the development of cognitive science. Though the field comes to be explicitly defined as such only in the second era, cybernetics is a crucial precursor to cognitive science and an integral part of its history. There are many lines of direct filiation between researchers, and this connection has been most fully elaborated and defended by French cognitive scientist Jean-Pierre Dupuy, who calls cybernetics the ‘poorly loved parent’ of cognitive science (Dupuy, 2000). For the first time with cybernetics, “the human brain was no longer comparable to the mechanism of an automaton or a clock; rather it *resembled*, and *functioned as, a computer*” (Changeux, 1997, p. 38). All three periods are thus linked together by their shared vision – by turns a metaphor, a heuristic, and a homology – of cognition as computation. They are split, however, by a series of interconnected, ‘intercalated’ transitions in both the hardware of the researchers’ computing machines and their metaphors of mental ‘software.’

The history of cognitive science is especially significant to my broader inquiry into media effects because its distinctiveness lies precisely in its status as deeply mediated form of psychology. It is an approach to psychology which aspires to transcend its status as a ‘human science’ and become a universal science of cognition: akin to McLuhan’s move in the analysis of media, it takes aim not at

psyche as the unique ‘content’ of biological brains, but something much more general and formal. In the era of cybernetics, this totalizing agenda was even more evident, centred though on concepts of communication, control, and feedback rather than cognition in particular. Throughout, this field is defined by a notion – first implicit, then explicit – of ‘multiple realizability.’ This is the vision that the formal properties of our interlinked mind-brain system may be usefully abstracted from their messy, evolved biological implementation and replicated in an idealized, mathematically designed model. These models were to be implemented and rendered self-acting within the electronic hardware of computing machines, operating in the first generation on continuous and then later on discrete quantities. Though the concept was only developed in the symbolic era, multiple realizability was equally central to some of the earliest work on cognition in the cybernetic tradition, particularly Warren McCulloch and Walter Pitts’ rendering of a ‘logical calculus’ in a simplified mathematical model of interconnected neurons (McCulloch & Pitts, 1943).⁵ For some cognitive scientists, these models were to be understood like any others, as useful abstractions for understanding some real phenomenon of greater complexity, rendering it more tangible, manipulable, and comprehensible. Hence all recognized that computers were poised to reconfigure both our understanding of cognition and our own cognitive processes. Another tendency reached its apex in the symbolic era, however, in the time of greatest optimism for true artificial intelligence; most famously espoused by Newell and Simon, this was the markedly more controversial view that these models were on their

⁵ As John von Neumann elaborated in his discussion of their 1943 article, “they did not want to get tied up in the physiological and chemical complexities of what a neuron really is. They used what is known in mathematics as the axiomatic method, stating a few simple postulates and not being concerned with how nature manages to achieve such a gadget... they believed that the extremely amputated, simplified, idealized object which they axiomatized possessed the *essential* traits of the neuron, and that all else are incidental complications, which in a first analysis are better forgotten.” (Von Neumann, 1966, pp. 43–44).

way to a true ‘emancipation,’ as cognitive agents which might soon outpace the capacities of their original human referents.⁶

⁶ As already discussed above, while I wish to interrogate this position, my agenda is not one of ‘anti-reductionism,’ the most common impetus for critiques of materialist cognitive science and AI research. Even more than neuroscience, which also has grounds to be skeptical of some ‘critical’ accounts directed against it (Fitzgerald, Matusall, Skewes, & Roepstorff, 2014), cognitive science complicates the story of reductionism with its notion of multiple realizability. Cognitive scientists adopt a range of positions on the ontological status of the human mind, ranging from eliminativism to panpsychism to simple agnosticism. True eliminative reductionism, the view that mind is a fictitious epiphenomenon of nerve firings without any causal efficacy, is a view rarely espoused in the literature (though it has prominent advocates, e.g. (Churchland, 1981). Methodological reductionism, as Fitzgerald et al. discuss, is perhaps both necessary and valuable, provided its limitations are recognized by researchers and in the public discourse around that research. Elucidating both the benefits and costs of reductionism, in my view, is a key role for critique in the sciences of mind and brain.

3.1 : Cybernetics: electromechanical brains and analog minds

A technoscientific metaphysics.

My account of the first period in this history is necessarily quite summary and selective as compared with the breadth of scientific endeavour gathered under the aegis of ‘cybernetics.’ Several more detailed histories of the main cybernetics group are available (most notably Heims, 1991, 1993), but even these do not aim to cover the full sweep of the movement. From the beginning, with Norbert Wiener’s *Cybernetics, or Control and Communication in the Animal and the Machine*, first published by an obscure French house in 1948—the year I chose to mark the start of the period—researchers have also set out to tell its history even as it was being made. Another useful account from the era, in equal parts history and theory, is Pierre de Latil’s *Thinking by Machine* (Latil, 1956). For my own rough notion that cybernetics was the most direct and important precursor to cognitive science, I found extensive support in the work of Jean-Pierre Dupuy (2000). My criteria of selection here are not based around particular institutions, nations, individuals or theories, but directed toward this argument. Cybernetics was as much about biology, psychology, and sociology, but I engage with these aspects inly insofar as they intersect with the cognitive. My intent is to show that the idea of mind as computation has its roots in cybernetics, and to engage with the devices that made this metaphor concrete: starting with the wartime technologies of fire control in anti-aircraft guns and bombs, and proceeding to their offspring in the ‘Homeostat’ and ‘Perceptron,’ two early electronic models of mind.

There is also a wealth of scholarship in STS on cybernetics (e.g. Galison, 1994; Haraway, 1990; Kline, 2009; Pickering, 2002, 2011). After defining the cybernetic approach and situating its most important human and technological actors, I take up these debates. While we cannot lose sight

of the military lineage and authoritarian connotations of this ‘science of control,’ equally we must attend to the cyberneticians’ attempts to envision different futures. The core of this chapter focuses more on the technological actors than the human ones, since the humans’ stories have been told in greater detail, and already structure even the most rigorously nonmodern accounts of the field (Pickering, 2011; on the “nonmodern,” see Latour, 1993b). Though the interests of those affiliated with cybernetics were far more heterogeneous, comprising a transdisciplinary endeavour with aspirations to span the natural and social sciences (Bowker, 1993), I emphasize its cognitive dimensions. Cybernetics is the most direct ‘parent’ of cognitive science, and conversely one of the most important unifying lines of inquiry in cybernetics was the reformulation of longstanding problems in psychology and the philosophy of mind through technological metaphors and analog machines. Of these machines, I select three principal examples—WWII weapons technology, William Ross Ashby’s Homeostat, and Frank Rosenblatt’s Perceptron.

While I place machines in the foreground, my account of their influence is of course not intended to be deterministic. Rather, I understand their role in much the way that cybernetics would have, as affecting the development of science through a dialectical feedback process equally involving theories, communities of scientists, and broader social concerns. I conclude this section, then, by considering some of the reflexive arguments of scientists about how cybernetics and its machines could be purposively deployed in service of particular visions for the good society. This ties in with an account of the later development of the ‘cybernetic organism’ concept, or the *cyborg*, in both its wild spacefaring (Clynes & Kline, 1960) and mundane computer-operating forms (Licklider, 1960).

As a first approximation, cybernetics refers to a loosely affiliated group of researchers centred around Norbert Wiener and a series of ten conferences on ‘Feedback Mechanisms and

Circular Causal Systems in Biology and the Social Sciences,’ sponsored by the Josiah Macy Jr. Foundation.⁷ Though only one academic philosopher participated in only a few of the Macy Conferences (Heims, 1993), and their focus was strongly empirical, the project was nonetheless also imbued with profound philosophical significance. Warren McCulloch, one of the primary actors in organizing the group, recalled Clerk Maxwell,

who wanted nothing more than to know the relation between thoughts and the molecular motions of the brain, [yet] cut short his query with the memorable phrase, ‘but does not the way to it lie through the very den of the metaphysician, strewn with the bones of former explorers and abhorred by every man of science?’ Let us peacefully answer the first half of his question ‘Yes,’ the second half ‘No,’ and then proceed serenely. (McCulloch, 1954, pp. 18–19)

This annexation of the metaphysicians’ territory through the ‘great heresy...of the knower as computing machine’ (*ibid.*) is integral to cognitive science and cybernetics. As Tara Abraham contends, “investigation of the mind-brain— and transcendence of the traditional dichotomy between the mind and brain,” was at the heart of the cybernetic project (T. H. Abraham, 2012, p. 559). It had a broader scope, but most participants agreed with McCulloch that the path to understanding all its other domains of interest ran through psychology, neuroscience, and computation. Thus Martin Heidegger saw cybernetics as aiming to replace metaphysics and philosophy in our technoscientific era (Heidegger, 1976), and Jean-Pierre Dupuy, with Heidegger in mind as much as Popper, calls it a ‘metaphysical research program’ (Dupuy, 2000)⁸.

⁷ The title of the conferences would be changed to simply ‘cybernetics’ following the sixth meeting, in 1949, following the publication of Wiener’s book in which he presents the newly-coined term. This was suggested by longtime editor of the conference proceedings, and a key figure in the ‘second wave’ of cybernetics, Heinz von Foerster, who recalled that the proposal was enthusiastically approved and Wiener, “deeply touched, leaves [the] room to hide his tears” (American Society for Cybernetics, n.d.). von Foerster’s ‘second wave’ of cybernetics research necessarily falls outside the scope of this project, though post-cognitivist cognitive science marks a return to ‘third-wave’ cybernetics (Maturana & Varela, 1980; Hayles, 1994).

⁸ Dupuy cites Heidegger as saying in the *Spiegel* interview that “Cybernetics is the metaphysics of the atomic age,” a quotation which has now percolated into other works and does poetically encapsulate his argument. However, it does not appear in any English translation of the article. The actual exchange in question is as follows: “*Heidegger*: The role which philosophy has played up to now has been taken over by the sciences... Philosophy dissolves into the individual sciences: psychology, logic, and political science.

Metaphysics, here, should not stand as a vague synonym for ‘mysticism’ or ‘nonscience.’ Heideggerian metaphysics, rather, is about uncovering a ‘primary being,’ a ground in relation to which all other beings ‘find their place and purpose.’ His later philosophy poses an opposition between *contemplative*, authentically philosophical thought, which ‘poses the question of meaning and of Being, understood as the sudden appearance of things, which escapes all attempts at grasping it,’ and *calculative* thought, which seeks to model, mathematize, and master the causal organization of the world for human control and use (Dupuy, 2000, p. 17). Calculative thought is the hallmark of technoscience for Heidegger, and though critiques of this formulation have been offered within STS (e.g. Latour, 2004b, p. 233), one can readily understand why Heidegger would see cybernetics as the height of metaphysics, seeking to supplant traditional doctrines of God and purpose. ‘Only a god can save us,’ he claims (Heidegger, 1976), but the ‘god’ with which we are left is cybernetics.⁹

Wiener of course recognized this, as did many others interested in ‘teleological mechanisms:’ this concept refers precisely to the cybernetic reformulation of purposiveness in mechanistic, mathematical terms. As Darwin sought to naturalistically account for the appearance of teleology in biological evolution, so did cybernetics in behaviour. Also like Darwin, Wiener was concerned with the ways in which his scientific doctrine would ‘impinge on religion,’ eventually writing a book on the subject (N. Wiener, 1964). And as McCulloch recognized, even the cybernetic ‘heresy’ of mind-

Interviewer: And now what or who takes the place of philosophy?

Heidegger: Cybernetics.” (Heidegger, 1976, p. 279)

⁹ As an alternative ‘overcoming’ of metaphysics, Heidegger poses his notion of ‘fundamental ontology,’ a contemplative thought which abides in the questioning of ‘Being’ as such and in general, without settling on a premature answer to the fundamental metaphysical ‘why’ question, be it a God or a mathematical world of forms. Neither cybernetics nor religion is philosophically satisfying on his view; but where the traditional mythology at least has the benefit of suggesting an infinite primary being which absolutely transcends human mastery, the new cybernetic metaphysics is directed precisely toward ‘steering,’ governing, and controlling the universe. This ties directly with his fears regarding technological exploitation of nature and ‘enframing’ (Heidegger, 1977). Digressing to the perpetual debates about Heidegger’s affiliation with National Socialism, it should be noted that these ‘ecological’ concerns in ‘The question concerning technology’ and elsewhere do not separate him from that movement, but rather ally him to ‘green,’ conservation-oriented elements within it, and within Hitler’s own thought (Snyder, 2015).

as-computation was not a wholly new idea. Histories of cognitive science and artificial intelligence often gesture back to Thomas Hobbes' equation of ratiocination with computation, but the roots of this idea extend back further still, to the very origins of philosophy. For cybernetics much as for Platonic metaphysics, the 'primary being' in relation to which all others are situated is a domain of mathematical abstraction. Its formal realm is just conceived using different tools. In place of the demonstrative, apodictic methods of geometry which were paradigmatic for the Greeks, cybernetics' world is conceived on the model of statistics and information theory. More importantly, cybernetics is afforded special access to this world, as it is made material and tangible through the new *technē* of electrical engineering.

Cybernetics thus begins with the analogy between humans and computing machines – or more properly, with the goal of understanding and modeling biological, mechanical, and electronic systems using a common mathematical language. (The notion of the 'cyborg,' of literally hybridizing the three, comes later.) It is indeed a metaphysical research programme in the Popperian sense, not a falsifiable theory but a framework for hypothesis-formation, untestable in itself. It “raises detailed problems in many fields, and it tells us what we would expect of an acceptable solution of these problems” (Popper, 1978). Equally, it has been observed that cybernetics proposes a distinct ontology (Galison, 1994). These philosophical dimensions of the field are worth disentangling. The Heideggerian conception is surprisingly congruous with the style of metaphysics proposed in analytic philosophy by David Lewis, as the inquiry into what *is*, fundamentally, and how these fundamental entities (inter)act to produce the phenomena we observe (Lewis, 1999). Where ontology is concerned with the full spectrum of existents, metaphysics seeks the primary substrate of being – for instance quantum particles or spacetime fields – from which all others emerge. Such questions as what it means for entities to be more fundamental than others, or what existence means

as opposed to nonexistence, are also properly metaphysical problems. Ontology, by contrast, encompasses the description and classification of the entire universe of things, most of which are aggregates, decomposable into more fundamental elements.¹⁰

The distinction is subtle but profound. Scientists, traditionally, have had far more to say about ontology than about metaphysics, and as Clerk Maxwell's symbolism suggests, have often seen this branch of philosophy as suspect. Wittgenstein and the logical positivists sought to overcome it altogether, carried along on the same philosophical trend as Heidegger, but far more enamored of science. I emphasize the metaphysical aspects of cybernetics, rejecting the common equation of the metaphysical and the non-scientific. This discipline begins with the analogies between a human pilot, human anti-aircraft gun operator, and the machines that bind them together in a tightly coupled system. This 'ontology of the enemy' and its military context is ultimately a small part of the picture, however, as cybernetics takes aim at much more fundamental questions. As Wiener argues, rejecting the scientific similes of earlier eras:

The mechanical brain does not secrete thought 'as the liver does bile,' as the earlier materialists claimed, nor does it put it out in the form of energy, as the muscle puts out its activity. Information is information, not matter or energy. No materialism which does not admit this can survive at the present day. (N. Wiener, 1961, p. 132)

The most enduring contribution of cybernetics lies not in the specifics of its ontology but in this new metaphysics it proposes. In a doctrine which would eventually be extrapolated into physics as 'It from Bit' (Wheeler, 1990), cybernetics situates information as the most fundamental reality, and sees the universe as a vast network of relays, switches and other mediators engaged in its processing.

¹⁰ For considerable further detail on the explicitly metaphysical speculations of the four 'synthesizers' among the Macy participants – along with Wiener and Bateson, he includes in this group the enigmatic John Stroud, and Filmer Northrop, the sole academic philosopher at the conferences – see (Heims, 1993, pp. 248–272).

From ordnance to Perceptrons: thinking machines in cybernetics.

The earliest point of conception for cybernetics, like so many other scientific and technological developments of the twentieth century, lay in the interdisciplinary, task-focused research teams established during World War II. As Galison has analysed in detail, and as Wiener himself clearly explained, his impetus was that the “speed of the airplane had rendered obsolete all classical methods of the direction of fire,” and so “it was necessary to build into the control apparatus all the computations necessary” (N. Wiener, 1961, p. 5). He called his ambitious and ultimately unsuccessful proposed solution to this problem the ‘AA predictor’ (Galison, 1994, p. 229). Working with the engineer Julian Bigelow, Wiener sought to design an electro-mechanical system to “usurp a specifically human function,” by executing a complex computation, and thereby “forecasting the future” (N. Wiener, 1961, p. 6). This was the first machine for thought in cybernetics, interpreted both as a ‘thing to think with’ (Turkle, 2007), and a thinking thing in itself.

Though enemy aircraft were flying nearly as fast as the shells which anti-aircraft emplacements were firing, if the pilots followed a simple straight-line trajectory, the computation would be simple enough for any operator, human or mechanical: they would simply need to ‘lead’ the enemy plane by a sufficient fixed distance that the shells’ path would intersect with its course. The technology which was eventually implemented, Hendrik W. Bode’s Director, did just this, continuously recomputing the target’s velocity and trajectory based on the past ten seconds, and it was both simpler and more effective than Wiener and Bigelow’s system (Galison, 1994, p. 244). The Director itself embodied some of the feedback principles which would become the focus of cybernetics, much like the earlier Watt governor. Pilots don’t simply fly in straight lines, though, and what made Wiener’s proposal unique was that it sought to incorporate the propensities of both allied gunner and enemy pilot into its model. Thus the machine did not simply employ feedback to

usurp some truly banal human function, like the apocryphal tale of the lazy valve-operator boy who accidentally invented the Watt governor with a bit of string. Instead, it incorporated a mathematical model *of* the human operators and pilots within itself, ‘smoothing’ the fluctuations in the gunner’s tracking while predicting the enemy’s zig-zagging path based on each pilot’s past performances. In other words, it incorporated *cognitive* elements.

Wiener and Bigelow’s predictor only ever existed as a simulation, using spots of light on a wall. But even during the war, they began extrapolating their models of feedback into a wide range of biological and psychological processes. The predictor was interpreted as modelling not only “the mind of an inaccessible Axis opponent but of the Allied antiaircraft gunner as well, and then even more widely to include the vast array of human proprioceptive and electrophysiological feedback systems” (Galison, 1994, p. 230). Pursuing the broadest possible extension of these ideas, a small circle formed to include Wiener’s closest collaborator, Arturo Rosenblueth, and together they developed the ideas published as ‘Behavior, Purpose, and Teleology’ in 1943 (Rosenblueth, Wiener, & Bigelow, 1943), first presented to a Macy Foundation conference in 1942 which included Warren McCulloch (N. Wiener, 1961, p. 12). By this time, cybernetics had been “fairly born but not yet christened” (*ibid.*, p.14). The network expanded to include eminent computer scientist and master organizer John von Neumann, as well as the team responsible for ENIAC and EDVAC, the first stored-program computers in the United States. Though Wiener was the primary actor in its ‘christening,’ Warren McCulloch was most directly responsible for organizing the aforementioned Macy Conferences, the primary site for its development. While Wiener’s emphasis was on engineering and physiology, McCulloch and Frank Fremont-Smith, a director of the Macy

Foundation, “rightly saw the psychological and sociological implications of the subject,” and so they “co-opted into the group” a number of leading figures in these fields (N. Wiener, 1961, p. 18).



Cybernetic Séance - New York City, 1947

From Left to Right : Ralph W. Gerard (Neurophysiologist), John von Neumann (Mathematician), Heinz von Foerster (Electrical Engineer), Lawrence K. Frank (Social Scientist), Norbert Wiener (Mathematician), Heinrich Klüver (Psychologist), Molly Harrower (Psychologist).
 Front (missing from view): Rafael Lorente De No (Neurophysiologist), Margaret Mead (Anthropologist), Kurt Lewin (Psychologist), Warren S. McCulloch (Neuropsychiatrist), Paul F. Lazarsfeld (Sociologist), Arturo Rosenblueth (Physiologist), Gregory Bateson (Anthropologist).
 Observers (missing from view): Frank Fremont-Smith (Medical Director of the Macy Foundation), Julian Bigelow (Computer Engineer), Walter Pitts (Mathematician), George Evelyn Hutchinson (Ecologist), Leonard J. Savage (Mathematician), Henry Brosin (Psychiatrist), Theodore Schneirla (Comparative Psychologist), Hans Lukas Teuber (Psychologist), Gerhardt von Bonin (Neuroanatomist), Lawrence S. Kubie (Psychiatrist), Filmer S. C. Northrop (Philosopher), Alex Bavelas (Social Psychologist) and Donald Marquis (Psychologist).

Figure 3 - 'Cybernetic Séance,' manipulated composition by Suzanne Treister

The first meeting of the still clumsily-named conferences on ‘circular causal systems’ took place at the Beekman Hotel in New York City on March 8-9, 1946. Wiener, Bigelow, Rosenblueth, von Neumann, and Fremont-Smith were all present, with McCulloch chairing the meeting. Other notable participants included McCulloch’s collaborator Walter Pitts, the

psychoanalyst Lawrence Kubie, Gregory Bateson and Margaret Mead, as well as the sociologist and mass communications scholar Paul Lazarsfeld (who lasted only a few meetings before apparently losing interest). Though Wiener emphasized the search for analogies between physiology and mechanism, “it was not so much the life sciences that it set out to confront.” Rather, the “first target of cybernetics was the sciences of mind” (Dupuy, 2000), in both the psychology of the ‘normal’ human subject, and the psychiatric treatment of pathology. Though those trained as mathematicians were nearly as large a group, psychologists were the largest single disciplinary group represented at the ten conferences. Along with McCulloch, and Foundation administrators Fremont-Smith and Lawrence Frank, who were interested primarily in psychology and sociology, other participating psychologists included Wolfgang Köhler, Molly Harrower, Kurt Lewin, Alex Bavelas, Joseph Licklider, and ten others, representing Gestaltist, psychophysiological, and psychoanalytic approaches to the field. At the time, “between the ‘hard line’ credo of the Establishment behaviorists and the unbridled conjecturing of the Freudians, it was difficult to focus in a scientifically respectable way on the territory of human thought processes” (Gardner, 2008, p. 15), and the Foundation was particularly interested in helping to stake out a ‘respectable’ compromise between these alternatives.

The linkages between biology, brain, mind and computer would be essential to the formation of this compromise, though as Heims describes, the exchanges between psychoanalysts like Kubie and neurophysiologists like McCulloch did not always produce cordial compromise. McCulloch publicly derided Freudianism as a ‘delusion,’ while Kubie speculated that despite—or perhaps because of—their genius, both McCulloch and Pitts were likely afflicted with some neurotic instability (Heims, 1993, p. 137). The case seemed stronger in the latter case, and Pitts eventually broke off contact with most of the group, seeming to grow increasingly troubled in his later years. In

a mysterious episode which involved Wiener's wife, a family scandal, and career rivalries, personal tensions between McCulloch and the Wieners would eventually lead to a complete break between the two, contributing to a breakdown of the Macy collective (Conway & Siegelman, 2009, pp. 223–224; Heims, 1993, p. 138). Nevertheless, some common ground was found even between psychoanalysis and the new computational picture of mind, with Kubie—like Jacques Lacan in France—eventually reconceptualising Freud's 'energetic' concepts of the unconscious with information-processing models (Dupuy, 2000, p. 116).

Whilst cybernetics and cognitive science would develop into the most potent alternative to behaviourist and psychoanalytic psychology (and psychiatry), these options were not initially seen as exclusive. Wiener's group was sympathetic to keeping the focus on behaviour and stimulus-response patterns, provided that feedback processes were inserted into the picture (Rosenblueth et al., 1943). As computers became more sophisticated, and the potentially universal capacities of switching networks and relays began to sink in, this machinery between stimulus and response began to look more like a mind, and less like a 'black box.' This was not due to any intrinsic property of the machines, however. Even as the cyberneticians were building their mathematical and material models, there were behaviorist psychologists designing machines in support of their own vision: E.G. Boring "used the computer model of mind to argue *against* mentalism" (Cohen-Cole, 2014, p. 161). Cybernetics and cognitive science chose to direct its research program toward the modelling of adaptive, autonomous behaviour, whereas Boring envisioned a robotics more in line with behaviorist understandings of conditioning and stimulus-response mapping (Boring, 1946; Cohen-Cole, 2005, p. 184). Yet the difference was not so great, for all the subsequent discourse of a 'cognitive revolution.' Boring and Wiener were in complete agreement that computers could simulate nearly any human capacity, and that constructing such simulations was extremely valuable

for psychological research (Boring, 1946, p. 178). Cybernetics set out less to reintroduce mind into psychology than to open up the 'black box' of behaviour by way of its central analogy, accessing the mechanisms of our behaviour by studying and tinkering with the behaviour of machines.

The Macy Conferences were the public, formal element of a broader circle then forming around Wiener, and just as computation was central to research they carried out, the group's meetings became central to the wider dissemination of the metaphysical research program centred on computing. Joseph Licklider, the psychologist, expert organizer and 'heterogeneous engineer' (Law, 1987) who went on to play a central role in cognitive science well beyond the era of cybernetics, recounts that he was first led to consider the implications of the digital computer for psychology by the "tremendous intellectual ferment in Cambridge after World War II. Norbert Wiener ran a weekly circle of 40 or 50 people who got together. They would gather together and talk for a couple of hours. I was a faithful adherent to that" (Licklider, 1988, p. 13). Considerable scientific and engineering work on these issues was equally under way in Britain, and the 'Ratio Club' was the counterpart to Wiener's circle there, a dining-club formed by J.A.V. Bates of those who "had Wiener's ideas before Wiener's book appeared," and wished to "talk 'Cybernetics' occasionally with beer and full bellies" (Husbands, Holland, & Wheeler, 2008, p. 99). In the British group still more than the American, the brain sciences were a focus of interest, with William Ross Ashby, William Grey Walter, and Horace Barlow among the founding members. But interdisciplinary cross-fertilization was just as important there, both amongst the sciences and with important actors in the development of computing technology: the Macy Conferences had von Neumann, the Ratio Club had Alan Turing, who joined at its second meeting. Connections between the two groups were not as strong as those within them, but McCulloch was a guest of the club on two occasions, while some of its representatives like Ashby and Donald MacKay were guests at Macy conferences.

Already at the time of the conferences, it was well understood that the digital approach to computation was preferable in many respects (N. Wiener, 1961, p. 117). The fundamental difference between digital and analog computation lies in its representation of data: analog machines, like Vannevar Bush's differential analyzer, represent and manipulate data in terms of a continuous scale, with their overall accuracy "determined by the accuracy of the construction of the scale" (*ibid.*) and the internal measuring instruments, while digital machines represent data by selection from discrete possibilities. By properly configuring and clearly distinguishing between these possibilities, any desired degree of accuracy may be obtained from a digital machine given sufficient computing power.¹¹ Two older technologies illustrate the distinction between analog and digital quite well: Pascal's calculator, and the slide rule. The former represents each digit separately and the value of each by distinct positions of mechanical gearing, while the latter represents numbers by the positions along a measuring line. Traditional calculating machines since the time of Pascal operated on our base-10 number system, but early computer scientists recognized that a machine could be equally capable and far simpler with only the two discrete possibilities of the 'binary' system. Computations could thus be carried out by a network of electronic switches and relays, first vacuum tubes and now microscopic transistors, while also supporting a foundational cybernetic analogy between the 'all-or-nothing' character of digital switches and biological neurons.¹² Assemblages of both could be

¹¹ Allan Newell offers another useful summary of this distinction: Analog computers represent quantities "by continuous physical variables, such as current or voltage; they were fast, operated simultaneously, and had inherently limited accuracy." Digital computers by contrast represent quantities by "discrete states;" in their infancy they "were slow, operated serially, and had inherently unlimited accuracy" (Newell, 1982, p. 9). Turing also gives the useful qualification that "strictly speaking," there is no such thing as a discrete-state machine, since "everything really moves continuously," but that the premise of digital computing is to make it feasible for operators to ignore this fact, and treat switching components as "definitely on or definitely off" (Turing, 1950).

¹² As Wiener summarized just a few pages on from his discussion of digital and analogue, "it is a noteworthy fact that the human and animal nervous systems, which are known to be capable of the work of a computation system, contain elements which are ideally suited to tact as relays. These elements are the so-called *neurons* or nerve cells. While they show rather complicated properties under the influence of electrical currents, in their ordinary physiological action they conform very nearly to the 'all-or-none' principle" (N. Wiener, 1961, p. 120).

interpreted as implementing symbolic logic (McCulloch & Pitts, 1943), and as I discuss below, this metaphor is concretely implemented in the design first of the Perceptron, then later in simulated neural networks running as software on digital computers. In the era of cybernetics, however, programmable digital computers like ENIAC are physically massive, expensive, and their processing power was largely monopolized by direct military applications, like the Manhattan Project. They were integral to the 'Big Science' of the era, and to the speculative futures of cybernetics, but in terms of the concrete practices of the time, cheaper, more widely available analog technologies were central.

The question of cybernetics' relationship to Big Science and the military-industrial complex is a topic of active debate within science and technology studies. Originating as it did with the Second World War and weapons technology, there can be little doubt that cybernetics is deeply entangled with power and violence. Computing itself was an outgrowth of military codebreaking efforts. Later artificial intelligence and cognitive science research just as evidently follows along the same path, although as we shall see there have been occasional droughts in funding. Nonetheless, Donna Haraway famously invoked the cyborg as a mythic paradigm of feminist hybridity: a figure on the borderline between science fiction and social reality, "resolutely committed to partiality, irony, intimacy, and perversity," while "oppositional, utopian, and completely without innocence" (Haraway, 1990, p. 151).¹³ She was by no means unaware of the concept's military origins, but saw the cyborg as a creature that defied this patriarchal filiation. Other writers have both implicitly and explicitly rejected this rendering of cybernetics, however, with Bowker emphasizing its imperious universalism (Bowker, 1993), and Peter Galison arguing in 'the Ontology of the Enemy' that

¹³ The 'cybernetic organism,' to which I return below, is a creation that postdates the period of cybernetics proper, postulated originally in the form of a human being modified with technological prostheses and drugs to travel through space without a ship (Clynes & Kline, 1960).

Wiener's vision of a 'coldblooded, machine-like opponent' in his AA predictor work furnished the outlines of an ontology in which disorder and entropy become the universal enemy. It is of the 'utmost importance,' he contends, "not to seize uncritically the central metaphors of operational analysis, game theory, and cybernetics and make them our own while claiming all the while a new 'postmodern' periodization" (Galison, 1994, p. 260). By way of further contrast, Andrew Pickering almost does just that, selecting quite different figures than Wiener to structure his history of the field, and offering a subtly different argument that cybernetics enacts a *nonmodern* ontology. As proposed by Bruno Latour, this implies a rejection of the nature-culture divide without reduction to either pole, and cybernetics does seem to fit the bill in this respect, with Pickering able to offer a persuasive reading of the field as something lively and performative, a 'minor' and 'nomad science' (Deleuze & Guattari, 1987) forming rhizomatic alliances with counterculture, Eastern spirituality, radical psychiatry and myriad reimagings of mid-century power relations (Pickering, 2011).

Which is it, then? Cybernetics is a name which can be applied to such a diverse collection of researchers and projects that one can readily find examples supporting either reading. We may choose to emphasize its aspirations to become a universal science of control, and its enduring linkages with the military-industrial complex, organizational theory, and the apparatuses of power. Alternatively we may see it as a nomad science, a practice of wild tinkering which aimed to understand mind and society, not for more effective and authoritarian governance, but in pursuit of a more open and democratic world. The only *wrong* answer is to imagine the field as represented in its entirety by one overarching 'cybernetic vision.' Rather, despite seeking a shared vocabulary for information and control in biological, mechanical, and cognitive systems, the theories, ideologies and research practices of cybernetics remained pluralistic. What unity there was came from its tools, not its ideas - and even this was unity of a discontinuous, meandering sort.

The first crucial thinking machines for cybernetics were weapons technologies – not only Wiener and Bigelow’s AA predictor, but also the guided torpedo which furnished a model for rethinking purposeful behaviour in their landmark paper with Rosenblueth (Rosenblueth et al., 1943). There, they suggested that while all mechanisms *can* be seen as purposeful, this is not an especially productive way to think, and further distinctions should be drawn. Some, like the roulette wheel, are expressly designed to accomplish no particular end. Others, like the gun, have an extrinsic purpose – an aim in the literal sense - supplied by a human operator. The guided torpedo, by contrast, is an intrinsically purposeful machine (Rosenblueth et al., 1943, p. 20). While its ends are supplied in some sense by the designers, it pursues them autonomously by a feedback process, calculating its distance to the target and adjusting its course and acceleration continuously until they collide. The nascent intuition of cybernetics was that whether the guidance unit were an electromechanical system or a human pilot, they might usefully be described in the same terms.¹⁴ The devices, then, are ‘thick’ things (Bijker, 2007), furnishing perfect object lessons in the genealogy of cybernetics. These weapons technologies indicate the longstanding ties between cybernetics, military and war, and epitomize an era when indeed cybernetics, “that science-as-steersman, made an angel of control and a devil of disorder” (Galison, 1994). This particular cybernetic vision was a product of the Second World War, when Wiener along with so many other scientists were eager to put their skills to use in service of the war effort. It remained central to later developments in

¹⁴ W. Grey Walter in fact gives a fascinating description, no doubt intended to emphasize the independence of the British tradition in cybernetics, of how a group centred on Kenneth Craik went through much the same thought process driven by military technology: “When [Craik] was engaged on a war job for the Government, he came to get the help of our automatic analyser with some very complicated curves he had obtained, curves relating to the aiming errors of air gunners. Goal-seeking missiles were literally much in the air in those days; so, in our minds, were scanning mechanisms. Long before the home study was turned into a workshop, the two ideas, goal-seeking and scanning, had combined as the essential mechanical conception of a working model that would behave like a very simple animal” (Walter, 1953, p. 125). Though coloured by a certain patriotism, Walter’s recollection is well attested by others, and like many in the Ratio Club, he saw their group as deeply indebted to Craik, and British cybernetics as profoundly set back by his early death in 1945 from a cycling accident.

human-computer interaction, as chronicled in detail by Paul Edwards' study of 'closed-world discourse' in the Cold War era of nuclear fear and the Semi-Automatic Ground Environment (SAGE) early-warning system. Wiener himself, though, later came to regret many of the military applications of his ideas, and worry about their implications for labour and inequality as well.¹⁵

From its origins, the visions of cybernetics have equally been concerned with peace. As Dupuy and Heims frame them, the Macy Conferences are best understood as a work of bridge-building between neurobiology, behavioristic psychology, and psychoanalytic psychiatry. Its funders were driven by a "faith in the curative, liberating, and pacifying power of the human sciences" (Dupuy, 2000, p. 82), particularly psychiatry, contingent on their cybernetic reformation. The World Federation for Mental Health, a closely associated group which counted Frank Fremont-Smith and Margaret Mead among its founders, took as its motto the phrase "Since wars begin in the minds of men, it is in the minds of men that the defense of peace must be constructed." Hence we cannot lose sight of either vision: one oriented toward war, the other peace; one pursuing the 'science of control' in service of traditional political, military, and corporate hierarchies, the other seeking to sketch out 'another future,' as Pickering puts it, using electronic circuitry as a means of reimagining and restructuring relations amongst peoples in positive ways. And as emblematic of this countervailing vision, a distinct 'minor,' nomadic, tinkering technical practice, we have Ross Ashby's Homeostat, a hacked-together configuration of "ex-RAF bomb control switch gear kits" (*fig. 4*) supposedly assembled around the kitchen table. While it is central to the alternative origin story that Pickering rightly wants to tell, of a cybernetics organized around the brain sciences rather than

¹⁵ As Galison puts it, focusing on the military aspect: "Paradoxically, during the war Wiener had extended the cybernetic vision beyond its narrow applications because of the weakness of the AA predictor; now that he associated cybernetics with the power of cataclysmic weapons, he tried to push cybernetics away from the military arena because of its deadly efficacy" (Galison, 1994, p. 255). This led to tensions in his once-cordial relations with scientists who maintained active participation in nuclear weapons research, like von Neumann. I return to Wiener's postwar thought below (*1c*).

military technology, its material filiation derives just as directly from the latter as its conceptual focus does from the former.

In Ashby's notebooks, he clipped a *Daily Herald* article of 13 Dec. 1948. Entitled 'The Clicking Brain is Cleverer than Man's,' it made rather dramatic claims for his machine, in one of the earliest instances of media hype for the cybernetic mechanization of mind. The Homeostat was characterized as a "demonstration 'thinking machine' which may one day be developed into an 'artificial brain' more powerful than any human intellect and capable of tackling the world's political and economic problems" (Ashby papers, p.2555). The machine "looks like four car batteries, bristling with switches," whose purpose is to model biological homeostasis in an electromechanical system.¹⁶ It was intended to represent not the very simple kind of homeostasis we might say is present in any thermostat or other feedback mechanism, but something that Ashby called 'ultrastability,' whereby even when placed into a chaotic and disturbed state, it stabilizes itself through one of several possible pathways. Where the first generation of cybernetics may well have demonized noise, disorder, and randomness, Ashby elevates it to a core design principle. The individual homeostat units were composed of movable magnets mounted on pivots, placed within coils and attached to needles. The intensity of the current running through the coil would cause the magnet to rotate more or less quickly in the direction of the coils' winding, thus moving the needle; the needle in turn dipped into a trough of water (the clear Perspex arc, *fig. 8*) with electrodes at either




¹⁶ The notion of homeostasis was coined by Walter Cannon in the 1930s to denote biological processes which kept certain parameters, like body temperature, within a definite range. It dates back to Claude Bernard's 'fixity of the internal milieu,' a phrase with both Ashby and Grey Walter liked to cite. For both, their cybernetic models were really hobbies alongside their hospital work, where they were pioneers of EEG technology, but nevertheless it was a kind of tinkering strongly influenced by their professional labour.

end applying currents of +5 and -5 volts, respectively (*fig. 7*). This current was then conducted to a triode, which controlled the coil current.¹⁷

Thus it acts as a pure feedback machine: “the current determines the position of a part... whose position again influences the current” (Latil, 1956, p. 298). Even if the position of a magnet is disturbed, or the current reversed using a commutator switch on the machine, it will ‘seek’ an equilibrium position in the middle of the trough, halfway between the two electrodes. When the magnets are thus interfered with, the machine responds by clicking gently: “these clicks are its ‘thoughts,’” the *Herald* suggested, as it ‘decides’ which of its 390,625 potential configurations will remedy the disturbance (Ashby papers, p.2555). These supposed ‘thoughts’ are the ticks of the Homeostat’s ‘uniselector’ component cycling through 25 randomly-selected resistors, their values having been chosen from a published table of random numbers. Far from diminishing its resemblance to cognition, Ashby saw this process of ‘selection’ from an elementary random process as central to all adaptive behaviour (Ashby, 1960). The machine was built as a mediator for this idea, serving to make a conceptual model tangible and manipulable, while also producing surprising and unexpected results. It was a machine designed to be capable of surprising its designer. While not all members of the cybernetics group were receptive to this idea (Heims, 1993), Wiener himself described Ashby’s “unpurposeful random mechanism which seeks for its own purpose” as “one of the great philosophical contributions of the present day” (N. Wiener, 1967, p. 54).

¹⁷ The vacuum tube triode, first invented by Lee de Forest as the Audion in 1906, was the first widely used electronic amplifier, and a key element in all the devices of the cybernetic era. In many ways it is the material basis for the very idea of ‘control,’ inasmuch as it was the primary tool by which a low-power signal could be used to govern the flow of a larger current.

(Maisemore) 2 March '48
 Have completed my new four-unit machine. It has four ex RAF bomb control switch gear kits as base, with four cubical aluminium boxes,



Clockwise from top left:

- Figure 4 - Ashby's homeostat
- Figure 5 - Ashby with his four-homeostat setup
- Figure 6 - Diagram of single homeostat
- Figure 7 - Interconnections of four homeostats
- Figure 8 - detail of four-homeostat setup

The single homeostat was a rather unexciting sort of machine, but considerably more interesting and dynamic behaviours were produced when several identical homeostats were linked together. In a four-homeostat setup, the output currents from each unit simultaneously affected in turn the inputs of each other connected box. Ashby interpreted the interlinked homeostat setup as analogous to the coupling of organism with environment.¹⁸ One homeostat operating autonomously represented the brain, adaptively responding to the changing conditions set by manipulations of the settings on the other units, representing its environment. Upon manually setting the output currents using potentiometers, and potentially reversing some currents with the commutators, the ‘brain’ unit might find itself in one of two possible states. As Pickering outlines:

It might be, as Ashby would say, in a condition of stable equilibrium, meaning that the vane on top of the unit would come to rest in the middle of its range, corresponding by design to zero electrical output from the unit, and return there whenever any of the vanes on any of the units was given a small push. Or the unit might be unstable, meaning that its vane would be driven toward the limits of its range. In that event, the key bit of the homeostat’s circuitry would come into play. As the electrical output of the unit increased above some preset value, the relay would close and drive the uniselector to its next position. (Pickering, 2011, p. 104)

The uniselector would insert into the circuit the next in its fixed series of random resistor values, which in turn would alter the state of the whole system. It would thus go on stepping through these random values until a state of stable equilibrium was found, whatever the initial configuration of the manual components might be. The uniselectors in each unit could be selectively engaged or disengaged, increasing the number of possible configurations to the aforementioned maximum of

¹⁸ By his functionalist logic, which would carry over into later cognitive science as well, this interpretation of the Homeostat was no mere metaphor. In *Design for a Brain*, he outlines his view that “given an organism, its environment is defined as those variables whose changes affect the organism, and those variables which are changed by the organism’s behaviour. It is thus defined in a purely functional, not a material sense” (Ashby, 1960, p. 35). As he argued in his contribution to the discussion at the 1952 Macy Conference, “there can’t be a proper theory of the brain until there is a proper theory of the environment as well,” and the subject thus far has been “hampered by our not paying sufficient attention to the environmental half of the process” (quoted in Pickering, 2011, p. 105). No doubt he saw the Homeostat as an important contribution in this regard.

390,625 (25⁴). As such, the machine often took hours, or even days to settle into a stable state (Latil, 1956). The Homeostat setup thus enacted Ashby's own cybernetic vision, one which was oriented toward the *organism-environment* rather than the human-machine system. The interest in feedback, for him, stemmed from an interest in this dynamic process of adaptation between the two poles. It allowed for a functional description which broke down the anatomical and physical separation between system's parts (Ashby, 1960, p. 39), applicable equally to a kitten learning about fire or a bird chasing a butterfly, and to a cyclist and their bicycle (*ibid.*, p.36) or, for that matter, a 'man-airplane-radar-predictor-artillery' system (Galison, 1994, p. 252). Thus while significant, the difference between Ashby's and Wiener's cybernetics was one of emphasis, not of essence.

This simple model nevertheless produced dramatic claims in the British media of the time, with the Homeostat proclaimed as 'always right,' but 'selfish,' and proposals made that one day this clicking electromechanical brain could be set up to tackle the world's great social problems (Ashby papers, p.2555). It also drew the interest of Alan Turing, who saw considerable overlap with his own work on thinking machines. In a letter to Ashby, he recounted that "In working on the ACE [Automated Computing Engine, an early electronic computer] I am more interested in the possibility of producing models of the brain than in the practical applications to computing" (Turing, letter to Ashby, 19 Nov., 1946). Ashby's machine operates in a wholly analog fashion, using the continuous variation of voltage in the trough as its signalling current, and it is electromechanical, changing its configuration by physically inserting different resistors into the circuit with the uniselector. In this sense it is paradigmatic of the cybernetic approach to modeling the mind. Turing, however, saw that "although the brain may in fact operate by changing its neuron circuits by the growth of axons and dendrites," a program could nevertheless be constructed for a digital computer which simulated this process without needing to materially alter the machine's construction. For his Homeostat

experiments, and in order to develop them further in a more flexible fashion, Turing suggested to Ashby that he would be “well advised to take advantage of this principle, and do your experiments on the ACE, instead of building a special machine. I should be very glad to help you over this” (*ibid.*). The Homeostat instantiated a set of equations which could be just as readily reproduced on a general-purpose computer. Ashby did not follow Turing’s advice, however, and instead sought to build another electromechanical analog machine which he called DAMS: the ‘dispersive and multistable system,’ an assemblage of electronic valves and neon lamps acting as relays, which was intended to be capable of taking on adaptive configurations as a system of many ‘ultrastable’ units like the homeostat (*fig. 7*).

Ashby insisted that building a true electronic brain now demanded little more than time and effort, and so he set to it, going ‘great guns’ on the new machine by July of 1950 (Ashby journal, p. 2953). Still seeing this as a diverting hobby in contrast with his clinical work, where he continued forging ahead with ECT, he built a small version on a shoestring budget. He concluded that a new principle was needed in the design for this machine, whereby it was in effect designed as little as possible, being

constructed from a set of elementary components analogous to neurons—in this case, neon lamps which insulated the flow of current below a certain threshold, above which they would switch over




Figure 9 – DAMS

and become conductors—which were then ‘assembled literally at random’ (*ibid.*). Theorizing it was better to ‘throw the design to the wind, and trust’ (p.2950), Ashby thus “built a semi-random contraption with 100 double triodes and watched it for two years before admitting defeat in the face of its incomprehensibly complex behavior” (Conant 1974, p.4). The unintended consequence of this aleatoric approach to simulating adaptive behaviour was that the simulation’s behaviour was no longer entirely comprehensible to its creator, and so its usefulness as a model was compromised. At the same time, it gave no indication that it might accomplish the alternative practical goal of simply accomplishing some useful ‘cognitive’ tasks, irrespective of whether it might shed light on how human brains do the same.

Pierre de Latil recounts that by 1953 Ashby already considered “that the present DAMS machine is too simple and is planning another with even more complex action,” but that “its construction would be an extremely complex undertaking and is not to be envisaged for the present” (de Latil, p. 310). One conclusion he reached was that it might be possible to simulate neurosis in this machine by what Gregory Bateson would later call a double-bind: “arranging the environment so that it affects two (or more) essential variables so that it is impossible that both should be satisfied” (journals, p.3463). Likewise, he saw suggestive analogies between the process of resetting the machine’s switching circuits and his clinical practice of ECT. Though Ashby recorded that “as usual, the designing forced a number of purely psychological problems into the open” (journals, p. 2967), DAMS was never really a success in his eyes, and references to it were purged from later editions of *Design for a Brain* (for a more detailed account see Pickering, 2011, pp. 124–132). Thereafter, Ashby took a position at the University of Illinois, where he continued to construct smaller demonstration machines according to the same principles as DAMS for pedagogical purposes, and further explored his own distinctive version of the cybernetic research

programme, at its core a vision that random processes and evolutionary selection constituted the foundations of all adaptive behaviour. Later on, cognitive science would return to this program after it had languished in obscurity for some years. Ashby's mediating models were certainly influential in the thinking of many contemporaries, particularly amongst the Ratio Club, but how to interpret these hacked-together assemblages of military-industrial castoffs was uncertain.

William Grey Walter, another noted British cyberneticist and one-time guest of the Macy Conferences—who likewise went by his middle name—once apologized for missing a meeting of the Ratio due to the “delivery of a male homeostat which I was anxious to get into commission as soon as possible.” By this, he meant the birth of his son (Husbands et al., 2008, p. 114). He also poked fun at Ashby's machines as devices ‘designed to do nothing,’ less a remark about their lack of



Figure 10 - Grey Walter's tortoise and 'cyborg family'

significance than about their immobility. Walter, for his part, also constructed machines “cobbled together from wartime surplus components and scrap materials” (Holland, 2003, p. 2093) which he called ‘tortoises’ – pictured together with the aforementioned ‘male homeostat’ in a portrait of the ‘cyborg family’ (Pickering, 2011, p. 2; original in Latil, 1956, p. 34). Walter even gave their ‘species’ a whimsical taxonomic name, using a Latin phrase that effectively means ‘machine to think with:’ *machina speculatrix*. Like Ashby’s Homeostat, they were comprised of simple feedback circuits intended to model adaptation, but by virtue of their being mounted on wheels and able to roll about the room, they were considerably more intuitive as models of *behavior*. In the parlance of robotics, they were no longer mere circuitry but machines with ‘effectors,’ components allowing them to act in and on the world. The tortoises were wired up with three wheels, two motors, a lamp, and a photocell light detector. Their default state was set to randomly steer around the room in a wandering pattern, cutting off the rotation of the front wheel when another light was detected so that it would head toward the source. When illumination reached a certain threshold, it would switch back into a wandering mode, unless the batteries were low, which would attract them to their illuminated base stations for recharging (descriptions derived from Latil, 1956; Pickering, 2011).

These simple rules, and the considerable role they allowed for random variation, would give rise to an array of surprisingly lifelike behaviours, again particularly when two units were combined together. The two tortoises would approach one another and then back away in a continuing ‘dance,’ while they would also present a similar ‘narcissistic’ response to their own reflection in a mirror. Although these mechanisms were simpler still, their mobility and Grey Walter’s public characterizations of them meant they were more readily domesticated, and viewed as behaving in a ‘lifelike’ manner. Though Ashby seemed to relish his media appearances as well, Walter was apparently a great showman, and appeared several times in print and once on a BBC newsreel

demonstrating his machine, footage which still survives online but seems to have provoked the ire of some others in cybernetic and neurophysiological circles (Holland, 2003). de Latil shows how the robots were presented as part of the family, alongside the young boy whom Walter viewed no doubt only partially in jest as sharing fundamental characteristics with cybernetic machines (Figure 10). Walter also encoded a distinct conception of gender roles into his machines. The romantic connotations of having them named ‘Elmer’ and ‘Elsie,’ while they were fated to a continual dance of attraction and repulsion, accounted for much of their appeal. Elmer had a shell of dull bakelite plastic, while Elsie was clad in red Perspex, and Elsie was programmed with a narrower tolerance for light values, causing her to dart about more ‘neurotically’ while Elmer was more ‘relaxed and sedentary’ (Hayward, 2001, p. 623).¹⁹

These tortoises have quite direct contemporary descendants in the iRobot ‘Roomba’ vacuum designed by Rodney Brooks, whose nonrepresentational approach to robotics I discuss later in this section. Walter also produced some of the more interesting speculations about the wider social consequences of cybernetic technology, to which I return at the close of this section. First, though, I wish to connect this history of cybernetics with the rise of the first functioning and widely available connectionist model, Frank Rosenblatt’s Perceptron. This machine was created at the Cornell Aeronautics Laboratory by a group with minimal direct connection to either of the primary cybernetics groups. Yet it represents the first practical success of the program outlined by Ashby, and its transition from kitchen-table hobby to a mainstay of cognitive science – though its fortunes will vary over the years. Equally with the Perceptron the material technology comes full circle, from

¹⁹ On the point of neuroticism, Walter also constructed a machine intended to demonstrate conditioning called the Conditioned Reflex Analogue (CORA), which he first placed in a mobile device like the tortoise, but then made into a benchtop classroom model. Pickering shows that for both Walter and Ashby, though their model-building was mainly a hobby, they often drew speculative connections with their work on EEG and ECT, and Walter tried to reproduce different sorts of pathological behaviour with CORA (Pickering, 2011, pp. 64–71).

the amateur assembly of military-surplus components, to the development of purpose-built machines for military use (though at this point, only partially and without great success).

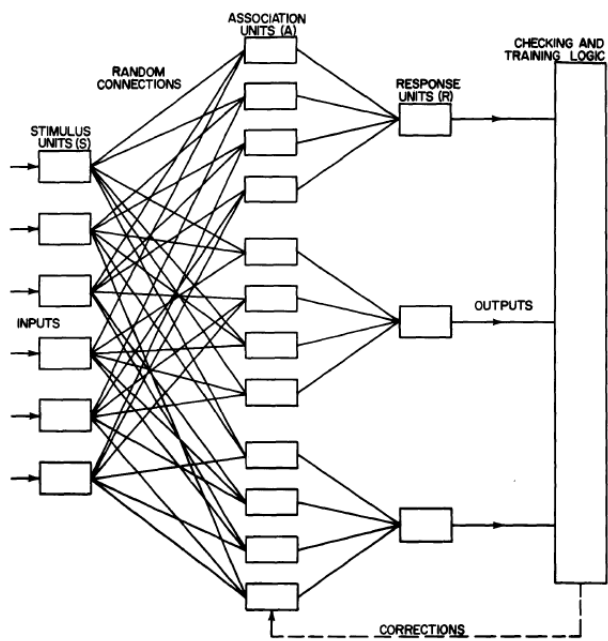


Figure 11 – Perceptron block diagram

When first proposing his model, Rosenblatt clearly acknowledged that Ashby's work was one of the most important precursors to his own, ranking Ashby as one of the few (along with John von Neumann) who concerned themselves with the functioning of 'imperfect neural networks' with 'many random connections' (Rosenblatt, 1958a, p. 388).²⁰ This paper was the advent of 'connectionism' as an approach to the mechanization of mind, and particularly of

memory. Rosenblatt coined this term to capture what he saw as an important theoretical opposition: between a view of memory as encoded in representational 'traces' that can be mapped individually to sensory patterns, "much as we might develop a photographic negative, or translate the pattern of charges in the 'memory' of a digital computer," and another, deriving from British empiricism, which views images as never really 'recorded' at all in this sense, but deriving from the signalling connections or pathways in the nervous system, viewed as an 'intricate switching network' (Rosenblatt, 1958a, p. 387). The Perceptron was a demonstration machine which he presented in support of this latter connectionist view, in an effort to concretize and experiment upon this

²⁰ The other figures to whom he acknowledges a substantial debt are McCulloch and Pitts, as well as Donald Hebb, the neurophysiologist who proposed that nerve cells which 'fire together, wire together,' and thus furnished additional biological plausibility for the Perceptron approach.

conception of the mind-as-switching-network. What was once a mere technological metaphor, Rosenblatt turned into a metaphorically-charged technology: both enacting a particular philosophy of mind, and designing a commercial product. The analogy between the Perceptron and biological systems was foundational to its design, and he noted that it should be ‘readily apparent’ to anyone who read his original proposal (*ibid.*).

By this point, digital computers were entering wider use within the military, proving useful for some kinds of tasks demanding high levels of mathematical accuracy, but less so for real-world problems like optical character recognition, or distinguishing moving targets in sonar. The United States Air Force had already attempted




Figure 12 - Organization of the Mark 1 Perceptron, from operator manual

without much success to put together an assemblage of Ashby’s Homeostats, in a project called Jenny, about which little information survives, but which was likely intended for similar tasks (Pickering, 2011); the Naval Ordnance Laboratory supported Rosenblatt’s work at Cornell with the intention of putting it to use in such ‘data reduction’ problems where traditional sequential methods were poorly suited (Pryor, 1961). Rosenblatt and his team had the same idea as Turing suggested to Ashby, however, and their machine was first run as a simulation on an IBM 704 computer, translating the fundamental principles of this adaptive feedback mechanism (*fig. 10*) into three pages of code in the ‘SHARE’ programming language (Pryor, 1961, pp. 35–39). Their now-declassified report concluded that the device, while impractical for usage on individual ships, would be a useful

tool for laboratory research on large bodies of sonar data, which would then inform improvements in the design of simpler ‘non-adaptive’ target classifiers for wider use (Pryor, 1961, p. 18). Although the Perceptron was an analog machine – based on continuous quantities, not to mention a fundamental analogy with biology – this distinction is by no means absolute, inasmuch as it first existed in the form of a mathematical model, and then a software simulation on a digital machine.

Subsequently, however, Rosenblatt’s work did lead to the production of a series of purpose-built Perceptron machines, starting with the ‘Mark 1,’ which were intended specifically for optical recognition (*fig. 11*). The Mark 1 implemented a 20x20 array of photoresistors in the film plane of a modified view camera, which would produce differing electrical signals depending on the contents of the scene, and could also be mounted in a box with a film projector for direct, automatic presentation of pre-photographed stimuli (Hay, Lynch, Smith, & Murray, 1960, p. 52). The varying electrical signals from these photoresistors were fed through ‘retinal’ circuits in a network of random connections leading to ‘association units.’ Though the machines have changed greatly, and once again exist principally as software simulations, the basic functioning of these units has remained the same in neural networks to this day. The individual elements are analogous to neurons, operating as relays with distinct activation ‘weights’ that are tuned through some adaptive process. Hence just as Ashby envisioned, they are designed in only a very elementary fashion, incorporating randomness in a significant way, and their knowledge is said to be *trained* rather than *programmed* in.

In the Mark I, each association unit contains a relay and a transistor amplifier, and its weights were encoded in potentiometers, adjusted by small DC motors in each unit (whose sound was in fact used as a cue in several of its operations, such as resetting the A-units: Hay et al., 1960, p.13). The unit’s input consists of “the algebraic sum of all the individual voltages on the connected input lines, and the relay will close only if this sum exceeds the A-unit threshold” (Hay et al., 1960,

p. 33). Each is thus connected to a random collection of optical elements, receiving a number of ‘activation’ signals from some photoresistors which are detecting light, and other ‘inhibitory’ signals from those which are not. If the sum of these activations exceeds its threshold, then it passes an output signal along the ‘many-to-one’ connections to the response units. These connections were variable by the operator using a plugboard, and the R-units also had their own threshold. If the sum total of inputs they received from the A-units exceeded that threshold, the R-unit would activate, illuminating a lamp on the machine. Different configurations of the plugboard could be used depending on the number of alternative possibilities the operator wanted to allow, and whether they were to be mutually exclusive; the machine could also be operated in either a ‘spontaneous association’ or ‘supervised learning’ mode. So if the intention were to train the Perceptron to recognize a few simple shapes, the operator would project a series of squares, circles, and triangles, either allowing the machine to repeatedly adjust its own A-unit thresholds, or deliberately using ‘reinforcement’ switches. Given a desired response, these would lead the A-units to produce it through a feedback process, repeatedly applying an error correction function to reduce the difference between the actual and desired output. In the former scenario, the machine would eventually settle into some stable pattern of responses, but they might not correspond with our categorization of the shapes; in the latter, it would be trained to give one of three mutually exclusive responses for squares, circles, and triangles, ideally being able to recognize them in images not part of the original training set. No one element of the machine could be said to contain *the* memory trace of the shapes, but the system as a whole constituted a distributed memory.²¹

²¹ These are the two key distinctions between the Perceptron approach and the serial, von Neumann architecture of digital computers: parallel processing, and the implementation of memory as distributed across the parallel processing units, rather than in a distinct memory ‘store’ of representations. This endures as a fundamental conceptual divide throughout the history of cognitive science. Though Perceptrons were first simulated on IBM computers, Rosenblatt argued that they represented a very different sort of model than the symbolic approach, which would conceive of the brain as directly analogous to a digital computer, and that such “models which conceive of the brain as a strictly digital,

The Mark 1 and its descendants were quite successful in some of these optical-recognition tasks, particularly when the stimuli were presented in a direct and clear fashion. They made concrete a particular formulation of the cognition-as-computation metaphor. But the same challenges observed by the early cyberneticists remained, and responding correctly in the face of noise and distortion proved elusive. Nonetheless like Wiener, Ashby and Walter, Rosenblatt had a certain showmanship about him. Another researcher called him a ‘real medicine man,’ and a ‘press-agent’s dream’ (McCorduck, 2004, p. 87), indicating some of the bitterness from the community which would later be reflected in the demise of the Perceptron approach. With the blessing of the Office of Naval Research, he gave interviews for a *New York Times* article which proclaimed that his machine would one day “be able to walk, talk, see, write, reproduce itself, and be conscious of its existence” (“NEW NAVY DEVICE LEARNS BY DOING; Psychologist Shows Embryo of Computer Designed to Read and Grow Wiser,” 1958). At the time, its practical successes were modest. In both its software simulation and Mark 1 implementations, it was capable of successfully discriminating between a few specific letters and shapes, or horizontal and vertical bars (Rosenblatt, 1960). It seemed to furnish an effective model for human perception, and a potentially useful device for pattern recognition tasks. It was at once a milestone in artificial intelligence research, and the culmination of the cybernetic approach to modeling the mind in electronic media. Though it could be translated with some precision into a sequential program for a digital computer, it was ideally conceived as analog, operating on continuously variable quantities encoded electromechanically, and it functioned adaptively, through a learning process based on feedback. How literally was the

Boolean algebra device, always involve either an impossibly large number of discrete elements, or else a precision of the ‘wiring diagram’ and synchronization of the system which is quite unlike the conditions observed in a biological nervous system” (Rosenblatt, 1958b, p. 422).

concept of ‘learning’ to be taken here? Rosenblatt contended that it was quite literal, an important step toward the construction of biologically plausible models for cognition.

There were many who agreed with him. As Marvin Minsky and Seymour Papert phrased it, ‘Rosenblatt’s schemes’ attracted “perhaps as many as a hundred groups, large and small, experimenting with the model either as a ‘learning machine’ or in the guise of ‘adaptive’ or ‘self-organizing’ networks or ‘automatic control’ systems” (Minsky & Papert, 1969, p. 19). Minsky and Papert were by no means pleased about this situation, however, and in their book – a more collegial, toned-down presentation of work that had been circulating as a polemical typescript – they set out to offer mathematical proof that labelling these machines as ‘learning’ was nothing more than a poor metaphor. At the same time, they were distinguishing themselves as founding members of a new approach to the research program of mind-as-computation, one which was more directly based on




Figure 13 - Cover of Minsky and Papert’s *Perceptrons*

digital computing than on biological analogies, and shifted focus away from the cybernetic watchwords of adaptation, self-organization, and control.

As the ‘official history’ of this controversy goes, Minsky and Papert’s *Perceptrons* showed that Rosenblatt’s simple neural networks, with a single layer of ‘association units,’ were unable to solve the crucial ‘exclusive-or’ (XOR) function, and so were destined for uselessness as models of cognition or as intelligent systems. Fitting the optical-recognition function of these early Perceptrons (later models would be applied to other problems such as speech recognition), the

book was in fact about ‘computational geometry,’ and they gave an illustration of their conclusions right on the cover. These two orange figures are clearly made up of two unconnected shapes, a fact which can be perceived by most human viewers. Minsky and Papert offered a mathematical proof, however, mediated by copious diagrams, that a Perceptron could never successfully determine the connectedness of such figures. This, again, is taken by many histories of AI and cognitive science to have been a definitive ‘falsification,’ or as close as one can achieve to such within this kind of technoscientific enterprise. All parties to the debate, though, noted that it was also difficult for humans to visually determine the connectedness of shapes like the internal green ones on the cover, at least without the mediation of a finger tracing the lines. Perceptron researchers suggested that this failing of their machine may be in fact testament to its biological plausibility. As Mikel Olazaran’s sociological analysis of the controversy contends, this book did not truly show that ‘progress was impossible’ in neural networks, and in fact researchers in the field were well aware of the XOR issue beforehand. Afterward some did continue to work in the area, simply being “displaced from artificial intelligence to other disciplines” (Olazaran, 1996, p. 613). The ‘falsification’ view, he argues, emerged as a result of the closure of the controversy for different reasons. It was because of ARPA’s subsequent defunding of neural network research and support of John McCarthy’s interconnected work on computer time-sharing and symbolic AI that Minsky and Papert’s book came to be seen as decisive, rather than the mathematical case driving the organizational decision (Olazaran, 1996, pp. 637–638; Edwards, 1997). Below, I look at one of the key human actors behind that shift, J.C.R. Licklider. Another technological factor was also at work, though. Historical accounts of AI often suggest that it was not until the 1980s and the work of the Parallel Distributed Processing group that the great discovery of multi-layer networks occurred. In fact, all parties to the initial controversy were well aware that in principle these networks could be constructed and would be capable of

much more. The missing ingredient was not the notion of ‘hidden layers,’ but good techniques for error correction and ‘training’ (Olazaran, 1996). Ultimately Minsky and Papert’s book was only the most public and formal presentation of an idea that had been solidifying in their research group – the ‘core set’ of the nascent symbolic AI. It was driven by a conviction that neural networks were the wrong approach, but as its later revival and the abiding controversy shows, this was less a solid fact than the expression of a basic social interest. They didn’t want the limited pool of funding for this research program to be directed toward this cybernetic, analog approach, and so against Rosenblatt’s protestations, they engineered a realignment.

Cybernetics’ speculative futures.

Perceptrons thereby serves as a marker for the end of cybernetics’ role in defining the mainstream of cognitive science. I selected three devices to draw out some of the manifold connections between tool-building and theorization in this period of cognitive science, using them to tell another looping story: from Wiener’s war work, to Ashby’s kitchen table strewn with military-surplus gadgetry, to the Cornell Aeronautical Laboratory and the Perceptron, this distinctive approach to the mechanization of mind came full circle. All of these researchers and tinkerers sought to make tangible the root metaphor of mind as information-processor, in the form of analog machines resembling the nervous system. The cybernetic movement was in part a speculative hobby taking flight from war work and materiel, following its own heterogeneous logics, but it also promised a set of highly desirable tools for military and administrative use.

It is worth noting that as American and British groups were each trying to construct their own model of science, technology, and social science for postwar liberal democracy, the Soviet

Union was pursuing two distinct programmes of their own with respect to cybernetics. In the first place a propaganda movement emerged, which has been convincingly presented as relatively 'bottom-up,' self-directed and self-sustaining in the Soviet popular press. (Gerovitch, 2001). It presented cybernetics as an evil tool of capitalist control, seeking to replace class-conscious workers with mindless machines, and promoting a misguided form of 'mechanism turned idealism,' in opposition to Marxist dialectical materialism and the work of homegrown materialistic psychologists like Pavlov. Their 'ideologically correct' alternative in the latter case was particularly dubious, in that it ignored "the fact that Pavlov's theory of conditional reflexes had been modeled on the telephone switchboard—a much more primitive mechanical metaphor than the cybernetic analogies with the servomechanism and the computer" (Gerovitch, 2001, p. 560), and most of these articles were based, at best, on brief summaries and extracts from Wiener's *Cybernetics*. In parallel with this public denunciation, started by a local editor but supported by Soviet authorities, the Soviet military leadership was also quietly recognizing that automated control and communication systems were a crucial strategic technology, in which they could ill afford to fall behind their Cold War rivals. The International Federation of Automatic Control (IFAC), one of many organizations which took up the concepts and problems of cybernetics under a different name, had its first conference in Moscow in 1960, and Edward Feigenbaum subsequently reported that there was considerable interest from the Soviet scientists, though his interactions with them were tightly controlled (Feigenbaum, 1961). He noted as well that there was concern about the "comparison between brains and machines" in an address by Norbert Wiener, likely informed by the ongoing press campaign against that specific element of Wiener's cybernetics. The "most vigorous applause" of the evening came when Wiener "stated his belief that the creativity of man would always find a higher level than

the creativity of the machine.” This “clearly gratified” the Russian scientists in attendance (Feigenbaum, 1961, p. 570).

The professional societies of cybernetics have splintered, but the IFAC, American Society for Cybernetics, and others endure. Cybernetics did indeed continue on through a ‘second’ and ‘third’ wave, according to some periodizations, centred respectively on Heinz von Foerster, then Francisco Varela and Humberto Maturana’s autopoiesis (Hayles, 1999). Its notions of feedback, information, control and self-organization are no longer the novel creations once so clearly identified with cybernetics, but they have percolated through a wide range of scientific disciplines – most notably genetics (Kay, 2000) – and into popular culture. By 1969, though, a transition already well under way seemed complete, and the cybernetic approach had ceased to be mainstream in AI and cognition research. It no longer commanded substantial military funding, though the total expenditure on all such projects was quite small in comparison with weapons research. Hype for its ‘clicking brains’ waned, and the collective interest of funding agencies, media, and the public shifted in a new direction. I discuss this rise of symbolic AI, its designers and its principal tools, in the next section; in the third and final section of this chapter, I contend that some currents of cybernetics find their way back in to the mainstream of cognitive science starting in the 1990s.

First, though, I consider some of the most interesting, lesser-studied contributions of cybernetics, when its attention turned to explicitly social and philosophical questions, and to speculation about what sorts of sociotechnical futures might come in its wake. The best-known of these is the concept of the cyborg, but that is only one in a broader genre of speculations about human-machine ‘symbiosis.’ The shorthand for ‘cybernetic organism’ has now eclipsed the historical origins of cybernetics itself. It does follow relatively directly from the notion that ‘control and communication’ were fundamentally similar processes in both animal and machine; the concept was

coined, however, in the context of speculation about how space travel might be pursued quite differently (Clynes & Kline, 1960). Manfred Clynes and Nathan Kline were an interesting pair from the Rockland State Hospital in Orangeburg, New York: the former was an autodidact Juilliard-trained concert pianist and electrical engineer who developed a specialized computer for the study of brain function, and the latter an early pioneer in psychopharmacology, best known for his work with reserpine. The cyborg concept, no doubt to their surprise, has proved their most widely-disseminated work. Rather than the conventional approach of constructing a bubble of the organism's natural environment in outer space, or in any life-hostile environment, they proposed that humans might respond to this not merely "technological" but "spiritual" challenge by taking an "active part" in our own "biological evolution" (*ibid.*, 26), and redesigning our bodies to survive in the environment of space.

There were a range of physiological difficulties in space travel that might better be solved this way, they argued. We might replace the function of the lungs with an artificial fuel cell, for instance, and the stomach and gastrointestinal tract by a direct nutrient-supply system. Also needed would be a "light-sensitive, chemically-regulated" system to adjust skin reflectance and absorption for the maintenance of body temperature. In keeping with their own interest in the brain and mind, however, they were also equally concerned with the psychological challenges of spaceflight, and so they proposed that an osmotic pump of the type recently developed for use in research mice could be implanted in the cyborg space-traveller, to administer continuous doses of drugs designed to promote wakefulness, improve perception and reaction time, or reduce disorientation and discomfort, "such as epinephrine, reserpine, digitalis, amphetamine, etc." (Clynes & Kline, 1960, p. 75). Though the concept would later be expanded to a matter of what humanity has 'always been' in our entanglements with technology (Haraway, 1990; Clark, 2003), the initial definition was much

more in line with the *Star Trek* image of the ‘Borg,’ as actual integrated biomechanical systems. More formally, they defined the cyborg as an “exogenously extended organizational complex functioning as an integrated homeostatic system unconsciously” (Clynes & Kline, 1960, p. 27). Just as we sweat on a hot day or shiver on a cold one, the engineered skin of the cyborg spacefarer would autonomously adjust its reflectance to keep them alive within a far greater range of temperatures. Their implanted drug pump would also be releasing drugs continuously, just as imperceptibly as – and in coordination with – the nervous system’s own endogenous hormones.

In a sense the cyborg was a far-flung speculative construct, a thought-experiment for how we might invert the premises of spaceflight in the far future. In another, though, it was clearly inspired by their own researches into drugs as behavioural technologies, and a real device, the Rose osmotic pump. This machine served as a springboard for their imagined unification of biology and technology, using autonomously-administered drugs not only to repair disordered, ‘pathological’ systems but alter and improve the functioning of ‘normal’ ones. And of course they ended their missive with that ubiquitous Cold War goad to *intereseement*: though some of these proposals “may appear fanciful, it should be noted that there are references in the Soviet technical literature to research in many of these same areas” (Clynes & Kline, 1960, p. 76). Though the cyborg in the specific sense they proposed was not a part of the first ‘wave’ of cybernetics – at that point the Macy conferences had ended and the core set of researchers had splintered – it certainly came to overshadow the work of that group in contemporary popular culture. And while neither the Americans nor the Soviets ever developed implantable devices for that purpose (at least to public knowledge), the core of Clynes and Kline’s proposal, namely the use of pharmaceuticals to push the limits of human capacities and modify cognition, has long remained an interest of the military and

intelligence services: most infamously in the CIA's MK-Ultra 'mind control' experiments with LSD, but also in the routine use of stimulants to ensure wakefulness on long missions.

Some of the mid-century Soviet scientists wary of Western cybernetics might have been surprised at its actual content, however, if not by its rhetoric. Certainly many Western writers did envision a world in which most of the labouring functions of humans – including an increasing share of 'intellectual' labour – were taken over by machines. Though no doubt a case could have been made that this was merely a bourgeois apologetics, they equally were concerned with maintaining a space for human creativity, as Wiener's remark to the Moscow congress implied. For one thing, despite the hype, they recognized that true intelligent machines which could match human capacities were a long way off. Equally, though, they were interested in outlining a positive vision for human flourishing in a world of ubiquitous computing. They contended that for the foreseeable future, we could accomplish far more by human-machine 'symbiosis,' employing the complementary skills of each, than either human reason or computing machines could acting autonomously. Though in the context of Cold War tensions this was in no way a question of direct influence, and in Wiener's writings he directly opposed his vision of society to the Marxist one, there were some interesting parallels between the cybernetic society and the ideal 'full communist' society proposed by Marx: characterized by a drastic reduction in labouring time, a reduction in the roles played by hierarchy and coercion in that labour which does persist, and a consequent turn of the masses toward leisure, intellectual pursuits, and consumption (Ollman, 1977). In many ways the ideas of cybernetics were quite similar, but drew inspiration from a more bourgeois, Keynesian vision of political economy. They also accorded a much more significant role to technological advancement than property reform in the reduction of human labour-hours.

W. Grey Walter presented an early vision of what the cognitive and social effects of ‘exteriorizing’ certain elements of our cognition might be. In good British style, as he did in his public characterizations of the tortoises, he employed a rhetoric of domesticity, arguing that “the exteriorization of tedious or controversial reasoning will no doubt have as profound an effect upon the brain and society as the introduction of skilled and respectful servants has on a humble household” (Walter, 1953, p. 194). It is only when these skilled new technological servants, with new modes of storing, retrieving, and processing information, “have done their work and retired unobtrusively to their quarters that the master brain can discover its own place and settle down to its proper work” (*ibid.*, p.234). Walter thus betrays a fondness for the traditional hierarchies of his society in the metaphors he uses to characterize the future co-evolution of brains and technologies; like the play of a country aristocrat as opposed to the wage-labour of the factory worker, the ‘proper work’ of the brain is equated by Walter with a kind of free and autonomous speculation. In keeping with his own modeling of the adaptive brain as a relatively simple goal-seeking mechanism, he imagined machines’ proper role as autonomously carrying out those mundane aspects of intellectual activities where goals could be clearly specified and feedback mechanisms constructed to achieve them, while the creative selection of projects and the difficult-to-formalize process of insight or ‘epiphany’ would remain the province of human brains, perhaps indefinitely.

In a sense Walter was transposing into a more specifically cognitive register the kind of futurism explored by John Maynard Keynes in his 1930 essay ‘Economic Possibilities for our Grandchildren,’ which considered the implications of increasing productivity. He concluded that within a century, all of humanity’s basic material needs – what he called *the* economic problem, the challenge of scarcity which had long been our overriding social concern – could be satisfied much more effectively with only a tiny fraction of the actual human labour currently expended (Keynes,

1963). Thus people's time would be given over to leisure and pleasure, leaving us 'only too glad to have small duties and tasks and routines.' This theory has recently given rise to some mournful musings about the failure of this world to materialize, and the rise instead of 'hyperemployment' within an attention economy (Bogost, 2013). Yet Keynes was quite open that while he saw this transition as already under way, its most visible first phase would be a great period of advancing *technological unemployment*, and a great adjustment of collective attitudes toward labour and accumulation of wealth would be necessary.

Norbert Wiener's meditations on the future of labour likewise sought to draw attention from the replacement of concrete, material labour by machines to the devaluing of cognitive labour. He wrote several times on this topic later in his career (N. Wiener, 1964, 1967), but already in the postwar introduction to *Cybernetics* this line of thinking was apparent:

Perhaps I may clarify the historical background of the present situation if I say that the first industrial revolution, the revolution of the "dark satanic mills" was the devaluation of the human arm by the competition of machinery. There is no rate of pay at which a United States pick-and-shovel laborer can live which is low enough to compete with the work of a steam shovel as an excavator. The modern industrial revolution is similarly bound to devalue the human brain, at least in its simpler and more routine decisions (N. Wiener, 1961, p. 27).

While again he preserved a role for the 'great mind' whose creativity – in complex and non-routine decisions – would remain indispensable for the foreseeable future, Wiener certainly did view this second industrial revolution as in many ways 'accomplished.' Whereas once 'computers' meant ranks of (typically female) human workers carrying out those aspects of mathematical calculation which could be clearly specified and routinized (*ie*, made algorithmic, akin to the rules of long division), now we knew that a sufficiently powerful Turing-equivalent computer could carry out any such calculation, with far greater speed and precision. Other whole sectors of labour, like the 'steno pool,' were on the decline. Wiener saw no end in principle to this path of development, and though he did acknowledge the increasing opportunities for employment in designing and maintaining this array of

machines, he looked toward a near future wherein “the average human being of mediocre attainments or less has nothing to sell that it is worth anyone’s money to buy” (N. Wiener, 1961, p. 28), neither in material nor mental labour. He contended that the answer was not to reject these technologies, as certain strands of anti-cybernetic thought in both the Soviet Union and the Western world implied (e.g. Ellul, 1964). Instead, he pointed toward a solution that many of his critics would likely have found congenial: “to have a society based on human values other than buying or selling” (N. Wiener, 1961, p. 28). Arriving at such a society would be a matter of social struggle, which might ideally play out ‘on the plane of ideas,’ or less pleasantly that of violent civil conflict.²²

Wiener saw his advocacy for labour interests and his pacifism as two sides of the same struggle, to ensure that the contributions of cybernetics on this ‘plane of ideas’ were of net benefit to humanity. As he saw it, he and his colleagues in the Macy Conferences and the Ratio Club had contributed “to the initiation of a new science which... embraces technical developments with great possibilities for good and for evil. We can only hand it over into the world that exists about us, and this is the world of Belsen and Hiroshima. We do not even have the choice of suppressing these new technical developments” (N. Wiener, 1961, p. 28). These twin horrors of Belsen and Hiroshima, those carried out by his wartime enemies and those carried out by his own nation, led Wiener from his conviction at the outset of the war that he needed to contribute as directly as possible to the military effort, to his still more fervent conviction at the end of his career that scientists needed to govern their research with a cautious eye toward its potential uses. There could be no more pure,

²² Here Wiener’s detailed writings may be contrasted with how Allan Newell characterizes the approach to this question of ‘Replacing versus Helping Humans’ in symbolic AI research: “An issue that surfaced about five years after the beginning of AI was whether the proper objective was to construct systems that replace humans entirely or to augment the human use of computers. The fundamentally ethical dimension of this issue is evident. Yet it was not overtly presented as an issue of social ethics, but rather as a matter of individual preference. An investigator would simply go on record one way or another, in the prefaces of his papers, so to speak. Yet there was often an overtone, if not of ethical superiority, of concordance with the highest ideals in the field” (Newell, 1982, p. 17).

objective and neutral science in his view, although this in no way diminished his estimation of its truth-value. Rather, when entwined with technology, and put to use in service of existing military and political powers, they could lead us not to Keynes' imagined state of 'economic bliss' but to a world of tyranny, civil unrest, or even the global supercatastrophe of a nuclear world war. His writings indicate the continuing pull of an *ethos* of scientific openness and free publication, but a growing wariness respecting the uses of his past research, and what should be undertaken in the future.

By late in his career Wiener sought to pursue only those projects, primarily psychological, which were furthest from military application in his estimation – though military interest in projects like the Homeostat and Perceptron may render this verdict a bit questionable – and he opened his *Human Use of Human Beings* with a letter to a fellow scientist employed by an aerospace corporation, who had requested a technical account of some wartime research Wiener had conducted. Again invoking Hiroshima and Nagasaki against the age-old 'comity of scholars' sharing information, he contended it was clear that "to provide scientific information is not a necessarily innocent act, and may entail the gravest consequences." The 'practical use of guided missiles' can only be the killing of enemies, and not the protection of American civilians – again an arguable claim – so although he is sure that 'with sufficient effort' his correspondent can find the information he seeks, he can only "rejoice at the fact" that his material was not readily available and protest '*pro forma*' in refusing to share it (N. Wiener, 1967, pp. xxvii–xxviii). Whatever the merits of his particular views, it should be clear that Wiener had a strong view of scientists' social responsibility, and viewed technological change as 'progress' only insofar as it hastened the end of military conflict and material want. The responsibility in respect of the former was to avoid research areas of direct military applicability; of

the latter, to advocate for social reforms to diminish the shock of an encroaching technological unemployment viewed as inevitable.²³

Wiener's increasingly outspoken pacifism on the one hand, and von Neumann's increasing involvement with classified weapons research on the other, acted to pull apart the Macy Conference group. By the last meeting, which had been relocated to Princeton for von Neumann's sake (though he was unable to attend in any case), the meetings were shedding participants, and they came to a rather ignominious ending. One assistant editor said this tenth meeting 'lacked content' and threatened to resign if required to assemble proceedings for it, eventually agreeing to the compromise that only papers and not discussion transcripts would be included (American Society for Cybernetics, n.d.). Tasked with summing up the conferences for the Foundation, McCulloch once again noted the heterogeneity of opinion and research interests among the participants, and said that while they didn't agree on much, the group felt that "we have learned to know one another a bit better, and to fight fair in our shirt sleeves" (*ibid.*). Again, I wish to emphasize that the real unity of the cybernetics group is not of theories but of tools: all were fascinated by the notion that we could mediate our understanding of mind through mathematically-based, biologically-analogous machines. Though the details of implementation were open for debate, this was the most enduring contribution of cybernetics to cognitive science. In defining itself as cognitive science, however, the field would renounce, at least temporarily, its parentage.

One comment and historical-sociological anecdote from Wiener's final book in fact gives a summary of two useful conceptual pairs for my interpretation of this history: "One of the most

²³ John von Neumann's thoughts on the matter furnish an interesting contrast, in support of his own continued weapons research. Reportedly he was more concerned about the potential impacts of computing technology and the "growing powers of machines" than of the nuclear bomb, fearing according to his wife that "What we are creating now is a monster whose influence is going to change history, provided there is any history left, yet it would be impossible not to see it through, not only for the military reasons, but it would also be unethical from the point of view of the scientists not to do what they know is feasible, no matter what terrible consequences it may have" (G. Dyson, 2012).

interesting social considerations in the sociology of invention is that of the interplay between the craftsman element and the purely scientific element” (N. Wiener, 1993, p. 61). The historical figures he adduces as having attained as ‘good a balance as was ever attempted’ in this regard were Michael Faraday and James Clerk Maxwell. Faraday the ‘laboratory boy’ developed his theory of electricity on the basis of a “language of images” and “figures of speech,” while Maxwell, “primarily a university man,” took these rough notions and translated them into “a sharp mathematical language” (*ibid.*). These are two closely related axes through which we might productively analyse any history of technoscience: the ‘craftsman’ as opposed to the ‘scientific’ elements, and the ‘language of images’ as opposed to that of mathematics. The generativity of cybernetics, through all its permutations inside, outside, and around the military-industrial complex, comes primarily from the interplay and translation between the poles of each. The autonomous computer model works to accomplish this with respect to metaphor and mathematics; standing in a comparable position to craftsman and scientist is the ‘organization man’ or heterogeneous engineer (Law, 1987). Wiener saw himself as an important figure of the Clerk Maxwell type, while others like Julian Bigelow were more wholly craftsmen of computation, and someone like Ashby or Walter stood somewhere in between. In alliance with the models, those researchers who proved expert at the tasks of heterogeneous engineering made themselves into crucial passage points, orchestrating the necessary strategic dispositions of funding, public interest, and academic institutions to create the enduring field of cognitive science. Wiener and McCulloch had been the most prominent of these actors for its first era, of analog and electromechanical models. Already taking their places for the next, however, were J.C.R. Licklider and Herbert Simon.

3.2 : Cognitivism: the ascent of digital representations

The emancipation of models and the return to minds.

Joseph Carl Robnett Licklider, known among close friends and colleagues as ‘Lick,’ was by training a psychologist; his work with the Harvard Psycho-Acoustic Laboratory and the Advanced Research Projects Agency has been well-chronicled (Edwards, 1997; Waldrop, 2001). Now he is more widely known as ‘computing’s Johnny Appleseed’ (Waldrop, 2001), for his role in disseminating the technology of digital computing around the United States, and also for his early role in proposing and directing funding towards ARPAnet, the precursor system to the Internet. In 1990 Robert W. Taylor observed—in an introduction to reprints of Licklider’s most significant papers on human-computer interaction—that “the least known of Lick’s accomplishments is perhaps his most significant” (in Licklider, 1990), framing the issue in more mundane institutional terms. Prior to his tenure, beginning in 1962, as head of the Information Processing Techniques Office at ARPA, there were no American universities offering Ph.Ds in computer science. By the end of his time there, he had directed funding towards the creation of four graduate programs in computer science which remain preeminent in the field, and figure centrally in the development of symbolic AI and cognitive science: at Carnegie Mellon, MIT, UC Berkeley, and Stanford (*ibid.*). He spoke once at the Macy Conferences, and was clearly familiar with their conceptual toolkit, arguing for the applicability of information theory in his own work on psychoacoustics (Dupuy, 2000, pp. 116–117). The spreading influence of Minsky and Papert’s attack on Perceptrons, and an interest in the related work of their MIT colleague John McCarthy on time-sharing and AI programming, led Licklider to direct financial support from ARPA toward this new approach, which in fact dated back to McCarthy, Minsky and Shannon’s summer research workshop at Dartmouth in 1956. This is the

year I choose to demarcate this period, recalled by Newell and Simon as the ‘critical year for the development of information processing psychology’ (Newell & Simon, 1972, p. 878), and generally recognized as the point of origin for both cognitive science and artificial intelligence. This section chronicles the history of this distinct era in cognitive science, though one overlapping with cybernetics both temporally and in some of the human actors involved. Here we reach the group which came to define the field as ‘cognitive science’ proper, and return where this chapter began, with a developing certainty that something was afoot whose implications spanned several disciplines. The equation of mind with electronic machines, first explored by cybernetics, would now be translated into digital media.

Licklider is a transitional figure, poised between the two eras. He had connections with the cybernetics group, but served in a similarly pivotal role to McCulloch and Wiener for a new set of researchers. A great many of those whom he employed at the Psycho-Acoustic Laboratory, and at centres funded through ARPA, became major figures in the ‘cognitive revolution.’²⁴ His own writings, like Wiener’s, help us to understand the context around the development of cognitive science as such a deeply mediated field. Like McLuhan, who described humanity as the ‘sex organs of the machine world,’ Licklider used the language of ‘symbiosis’ and ‘mechanically extended man’ to describe his vision for the future of human-machine interaction. While once he described it by the more staid notion of ‘partnership,’ eventually he came to the more suggestive analogy with the fig tree and its tiny pollinating wasps, *Blastophaga grossorum*. Though “conceding dominance in the distant future of cerebration to machines alone,” he rightly saw the main intellectual advances of the near term as most likely achieved through “men and computers working together in intimate association”

²⁴ Alongside many others, the coauthors of *Plans and the Structure of Behavior*, widely recognized as the landmark work of cognitive psychology—George A. Miller, Karl Pribram, and Eugene Galanter—started their careers at the PAL (Edwards, 1997, p. 212).

(Licklider, 1990, p. 3). Here we return once again to an economy of attention, which he contended was not paid out in a particularly efficient way by unaided humans. About 85 percent of ‘so-called thinking time’ was spent getting into a position to properly learn, reflect, and decide; much more time was spent retrieving and organizing information than properly understanding and making use of it (*ibid.*). Thus he imagined that even well before the development of true artificial intelligence, which might usurp our cognitive functions altogether, we might usefully delegate much of our cognition to machines. Our attention would be reserved for higher tasks. Though somewhat akin to Grey Walter’s speculations, Licklider’s imagined world is less playful and more practical.

Information-processing machines could be built to “convert hypotheses into testable models and then test the models against data,” to simulate mechanisms and models, and display results to human operators. Such a machine might transform data and plot graphs in multiple ways, and “convert static equations or logical statements into dynamic models so the human operator can examine their behavior. In general, it will carry out the routinizable, clerical operations that fill the intervals between decisions” (Licklider, 1990, p. 14). Licklider’s arguments stand as some of the most prescient regarding the future of computing and cognition, and serve equally well to describe the role played by digital media in this period of cognitive science. As Edwards contends, “The goals articulated in ‘Man-Computer Symbiosis’ became, almost without revision, the agenda of ARPA’s IPTO under Licklider: time-sharing, interactive computing, and artificial intelligence” (Edwards, 1997, p. 269)

Along with Vannevar Bush’s hypothetical ‘Memex’ system (Nyce & Kahn, 1991), Licklider also contributed substantially to the conceptual foundations of the Internet, and in his institutional role helped to steer ARPA toward constructing such a system. In another notable paper, “The Computer as a Communications Device” (1968), he contends that “Creative, interactive

communication requires a plastic or moldable medium that can be modeled, a dynamic medium in which premises will flow into consequences, and above all a common medium that can be contributed to and experimented with by all” (Licklider, 1990, p. 22), and that the computer will offer the technical foundation for such a shared, multidirectional, plastic medium. The problem with two-way telecommunication, he thought, was perhaps less to do with an inability to see the face of one’s interlocutors, and more with the lack of facility for creating and modifying shared external models, a mode of working with media that is essential to any real-world activity of collective cognition (Hutchins, 1996). Instead of pursuing videoconferencing exclusively, he saw the value in early forms of computer-mediated collaboration such as those developed by Douglas Engelbart, which already included many elements that have become the norm in such affairs today.²⁵ Licklider thought it might ultimately be more effective to situate the core set of researchers working on a problem at distributed sites, connected by computer and occasional personal visits, rather than together at the same institution in close physical proximity, for “the most creative people are often not the best team players, and there are not enough top positions in a single organization to keep them all happy” (Licklider, 1990, p. 29).

Although the material infrastructure for linking these together with computerized communication did not yet exist, as the first head of the Information Processing Techniques Office, Licklider set about creating the network of research sites which would make it possible. The logical map of the early ARPAnet bears his unmistakable stamp, showing the literal network connections constructed between the academic sites where computer science, cognitive science, and AI were

²⁵ In one model teleconference, each presenter, for instance, is described as having a ‘topical outline’ accessible with hyperlinks to important data; everyone in attendance has their own screen to ‘thumb through’ the speakers’ files and revise their own internal models more effectively (Licklider, 1990, pp. 22–29).

simultaneously being

developed.²⁶ He seems to

have been vindicated in

his intuition that

computers may increase

the effectiveness of

communication at a

distance, even

‘revolutionize’ the

process, but that they

remain incapable of truly matching human capacities or operating fully autonomously in the near

term. The development of cognitive science, like ARPAnet – and for that matter Ashby or Walter’s

cybernetic models – was a matter of humans acting together with machines, sometimes putting them

to use for definite ends, sometimes allowing the unexpected to emerge from their interaction. It was

a recursive, reiterative process: reshaping scientific thinking about thought by attempting to build

machines that think. And while the causal role should not be overstated, as much as the context of

cybernetics was shaped by wartime experience during the Second World War, and a dream of

subsequent peace (Heims, 1993), that of cognitive science and AI was by the Cold War. This played

out as an interest in both practical technologies for defense (Edwards, 1997), and conceptual

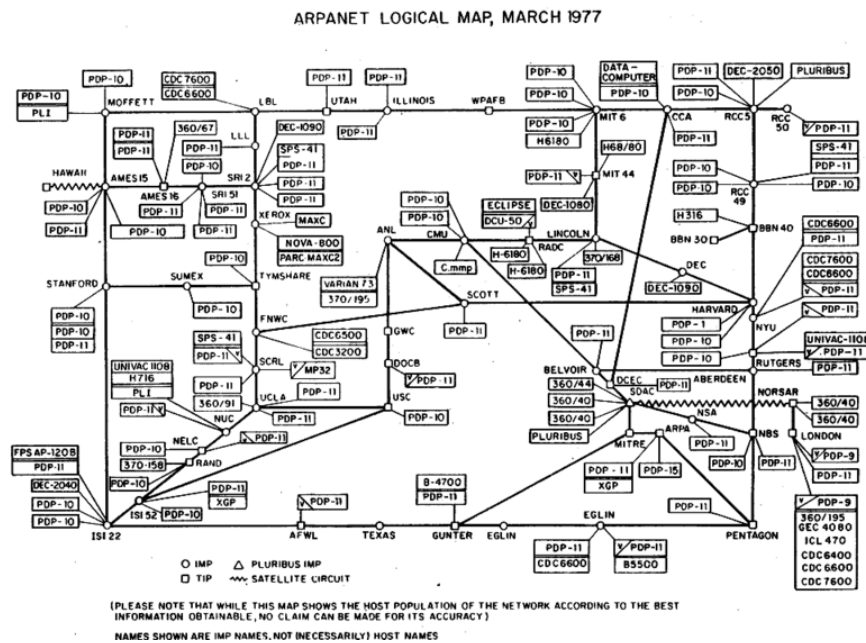


Figure 14 - ARPANET Logical map

²⁶ The other most prominent institution represented on the map, instrumental in the creation of the packet-switching technology which underpinned the network, was the consulting firm where Licklider worked prior to his tenure at ARPA: Bolt, Beranek and Newman. Like Licklider himself, the firm originally focused on acoustics and analog computing—known for its consultation on various concert halls, as well as analysis of the Nixon tapes—but eventually turned to digital computing more generally, with the rise of digital signal processing, becoming the first purchaser of a DEC PDP-1 machine and a key contractor for ARPA. Licklider recalls being careful to avoid conflicts of interest, and recusing himself from consideration of BBN bids while directing the IPTO (Licklider, 1988).

techniques for describing the different cognitive ‘types’ fostered under capitalism and socialism (Cohen-Cole, 2005, 2014).

Allen Newell’s historical account of AI gives a useful summary of the split between cybernetics and symbolic approaches. The ‘instant rise to prominence of cybernetics’ came with a solution to the dilemma of ‘mechanism versus teleology’ by Rosenbleuth, Wiener, and Bigelow. Newell recognizes this as a fundamental philosophical debate dating back to the Cartesian split of mind from matter, and says the cybernetic thesis that ‘purpose could be formed in machines by feedback’ was ‘universally perceived’ as groundbreaking, while the field’s later decline involved ‘no change of opinion on this issue’ (Newell, 1982, p. 6). The split derives, rather, from the intellectual opposition between ‘symbolic versus continuous systems’ and ‘psychology versus neurophysiology’ (*ibid.*, p.5).²⁷ Newell and the other core researchers of symbolic AI evidently laid their sympathies with the former, while cybernetics emphasized the latter. After “a certain amount of skirmishing,” by 1970, “the field of computers came to mean exclusively digital computers” (Newell, 1982, p. 9). The Fast Fourier Transform equations created the field of digital signal processing, which until that time had been the last major bastion of analog computing, and by the time he was writing in the 1980s, Newell noted that many young researchers “hardly know what analog computers are” (*ibid.*).²⁸ Cybernetics and the Perceptron tradition of cognitive modeling was oriented toward pattern

²⁷ Later he adds ‘Problem-Solving versus Recognition’ and ‘Serial versus Parallel’ to this cluster of issues around ‘AI versus Cybernetics,’ as well as the valuable caution that these clusters don’t really pick out one unitary ‘issue’ of opposition, and that scientists can always be found ‘aligned in non-standard patterns’ (Newell, 1982, pp. 35–36).

²⁸ He notes, moreover, a developing split between ‘pure’ computer science and artificial intelligence research with respect to the nature of representation in these new machines: the former “defined computers as machines that manipulated numbers. The great thing was, they said, that everything could be encoded into numbers, even instructions. In contrast, the scientists in AI saw computers as machines that manipulated symbols. The great thing was, they said, that everything could be encoded into symbols, even numbers” (Newell, 1982, p. 9). This subtle difference in outlook had the effect of isolating AI within computer science to some degree according to Newell, leaving it a part of the wider field but one with a special, somewhat suspect status – and more subject to cycles of media hype and backlash than more mundane practical uses of computing.

recognition, whereas the shift to symbolic AI was accompanied by a reorientation toward ‘problem-solving’ as the essential task of cognition.

On the material side, then, this new mode of cognitive science was characterized by digital technologies. The foundation of its mind-as-computation metaphor had already been laid by cybernetics, but computation was coming to mean something entirely different. This suggests that, contrary to the notion that metaphors exercise a truly determining influence on the questions and answers in science (Edge, 1974), they function through a dynamic interplay of the figurative and the concrete. Even where two groups have adopted the same generative technological metaphor, the class of actual systems which they adopt has an avowedly “profound influence on the course of a science” (Newell, 1982, p. 11). Here Newell outlines a kind of technologically-grounded Kuhnianism for his model-oriented science, arguing that “alternative theories which are expressed within the same class are comparable in many ways. But theories expressed in different classes of systems are almost totally incomparable. Even more, the scientist’s intuitions are tied strongly to the class of systems he adopts” (*ibid.*). The choice between digital and analog systems has a deep structuring effect on their intuitive characterization of the analogy between mind and computation, shaping the problems they target and the solutions they propose, but not in any unidirectional or deterministic causal fashion. Rather, as Newell goes on to acknowledge, these technical and conceptual choices are channeled in important ways through institutions. The ‘major historical effect’ of the analog-digital divide was a ‘rather complete separation’ of groups focusing on each, now ‘strongly institutionalized.’ The “continuous-system folk ended up in electrical engineering departments; the AI folk ended up in computer science departments,” and at the outset the former were “almost exclusively focused on hardware systems,” the latter on software (*ibid.*). This focus on software was of primary importance in the greatest leap of cognitive science toward mediation, the ‘emancipation

of the model' (Dupuy, 2000, p. 60) known as multiple realizability. Here we arrive at the strongest sense in which I contend that cognitive science is a *mediated* discipline, as the modelling of mind 'takes flight on its own wings' (*ibid.*), freed from its prior coupling with neurobiology.

Cybernetics did as much as any scientific programme had hitherto in constructing autonomous models of cognition – variously described in terms of adaptive behaviour, intelligence, brains, and minds. It saw these as something more than mere useful simulations, but as in some ways manifesting deeper homologies, in the processes of information transmission and feedback they represented. As I described at the outset of this chapter, starting with McCulloch and Pitts' translation between formal logic and abstract neural systems, on through Perceptrons, this was a central aspect of cybernetics, and the precursor to what would be defined as multiple realizability in cognitive science. As Newell puts it, "that human intelligence is guide and goad to the engineering of intelligent systems was clear" right from the start, yet this left open the question of whether inspiration would be drawn primarily from psychology or from neurophysiology, two disciplines which "speak with entirely separate, though not necessarily contradictory, voices" (Newell, 1982, p. 13). Cybernetics took the latter path, seeking to model cognition by replicating some key aspects of what we know to be the physical basis of our own intelligence—the brain and nervous system—in purpose-built analog or 'continuous' machines. Cognitive science and artificial intelligence, by contrast, sought to replicate more directly what we knew about the functioning of the *mind* through psychology, in software programs written for digital machines. Although both "treated thought as computation," cybernetics "continued to locate the agents of computation at the neuronal level," while assigning agency to an "autonomous level of mental representations, created what was to become the distinctive style of orthodox cognitive science, or 'cognitivism'" (Dupuy, 2000, p. 63).

‘Cognitivism,’ then, as Dupuy and others have used the term, implies the elevation of computing to a primary metaphysical status: the view that thought does not simply resemble computation, but that thinking *is* computation in the most literal sense. Cognition on this view is defined as the rule-governed manipulation of symbols which represent aspects of the world (Varela, Thompson, & Rosch, 1991), a single process with unified ‘laws’ which may be instantiated in biological or electronic ‘hardware.’ As Turing advised Ashby, biologically-inspired models could readily be simulated as programs on digital computers rather than requiring special devices, and so this was certainly a possible path of development even with the rise of digital machines. But cognitivism stripped out the intermediate layer of biological analogy and sought to directly reconstruct the mind. In the process, it made the scientific study of mind fully respectable once again, stripping away those vestiges of behaviourism which left so many cyberneticists reluctant to write the word.

Programming tools for cognition: Logic Theorist, General Problem Solver, and list processing.

Despite the foregoing argument, cognitivism was not a radical departure from the cybernetic program, but its greatest extension in the study of psychology. It took to heart what Wiener had already argued: that ‘information is information,’ deserving to be treated as a distinct object of scientific inquiry in its own right, and that digital computers were the best means of undertaking such inquiry (N. Wiener, 1961). At its origins, cognitivism was simply ‘information-processing psychology,’ and Joseph Licklider was directly influenced by Wiener and Rosenbleuth’s discussion circle in first considering the possibilities of digital rather than analog computing for psychology (Licklider, 1988). Herbert Simon, for his part, was reportedly not a great fan of Wiener’s systematization of cybernetics, but claimed to draw his inspiration from ‘classical servomechanism

theory' dating back earlier—again primarily to wartime research—and found Ashby's formulations more appealing, writing him in 1953 to say that *Design for a Brain* was “the most exciting book I have read in a decade” (in Heyck, 2008b, p. 54). Simon's writings contain many approving references to the work of Ashby and Walter, compared with few to the American cybernetics group. The break with cybernetics, which may not even have registered with some of the actors, was not a matter of choosing a different research problem or a different tool, but of the implications following from changes in the essential characteristics of the tools employed. Cognition and computation remained the central focus, and a two-way traffic from tools to theories and back continued to drive the research programme, but as digital machines took over from analog ones, the nature of the models changed. This change was not necessary but quite contingent.²⁹ It had much to do with the rise of a different social network of researchers, this time centred on Simon and Licklider, but in some ways as heterogeneous as the cybernetics group. They could even have kept at simulating neural networks on these new machines, but Minsky and Papert were important actors within this new group, and though as we saw their critique of Perceptrons did not wholly rule out improvements on the basic concept, it certainly signaled the group's antipathy to that approach. Instead, they developed a new type of modeling strategy to match the new capacities of stored-program digital computers.

²⁹ As Paul Edwards has persuasively argued, even the transition to digital machines was in no way an inevitable development. Though often told in the history of computing as the self-evident consequence of the superiority of digital technology, in the immediate postwar period of the 1950s, this was in no way obvious. Analog computing machines like Vannevar Bush's Differential Analyzer were the better-understood, more widely-used, and faster machines at the time, while digital machines were substantially slower, lacking the inherent 'real-time' operation of analog machines and instead operating at a set 'clock speed' of instruction cycles (Edwards, 1997). Their only clear advantage, as Wiener and others recognized, was the capacity to carry out calculations to any arbitrary degree of accuracy. The superiority of digital computers was not the cause of, but a fact that emerged as the consequence of an institutional shift in support of this approach, which I try to trace in brief within this section, and it was driven by a heterogeneous set of contextual motivations. Edwards focuses on the macrosocial factors that drove the rise of digital computing in general – their function as a support for 'closed-world discourse' – while I try to foreground the smaller-scale processes which drove the adoption of specific models and tools in cognitive science more narrowly, without losing sight of this broader picture.

To speak of the Homeostat or the tortoise of the cybernetic era as possessing a ‘representation’ of the world would be a stretch, but it is not impossible. Typically these are considered machines that operate without a detailed representation of the external world. The state of an individual Homeostat unit, or of the switching and signaling components in a tortoise, could be taken as embodying a representation of its environment, however; the differing weights of the nodes in a Perceptron more clearly embody a kind of distributed memory. In all these examples, infinitesimal variations in current imply representations of an ever-shifting, peculiar sort. A sufficiently complex system of this type, as Ashby indicated, might well cease to be fully intelligible to its designer in all of its workings. Cognitivism changes all this, granting a much more prominent role to a more intuitive sort of representation, embodied in discrete symbols. ‘Reprogramming’ the older analog machines, or simply carrying out the adaptive processes they were meant to model, required physical changes in their setup, evident in the clicking of the Homeostat or the hum of the Perceptron’s DC motors. The new machines, by contrast, were universal computers in the sense proposed by Turing, operating by the manipulation of simple abstract tokens (typically those of binary arithmetic) according to definite rules, and by reading out instructions (in those same tokens) for the modification of its own rules. These basic elements, later developed into the von Neumann architecture, give rise to machines which can compute anything computable: if it can be put in algorithmic form, spelled out into a series of step-by-step instructions, then such a computer can carry it out. The first such Turing-complete machine had been the ENIAC, built during the Second World War by John Mauchly and Presper Eckert to compute firing tables for the U.S. Army and put to use thereafter by von Neumann in designing the hydrogen bomb.³⁰ If Charles Babbage’s second

³⁰ Two other experimental computers were put into operation very shortly after ENIAC, the Institute for Advanced Study machine in Princeton, by a team led by Julian Bigelow and von Neumann, and the Manchester Small Scale Experimental Machine or ‘Baby’ in the U.K. Information about the IAS machine was widely shared, and the first

analytical engine had been completed, however, it would have been the first. Tasks like the calculation of ballistic trajectories, or logarithmic tables—Babbage’s original aim—are the paradigmatic computable problems.

Such problems come in the form of numerical representations, and they can be broken down into a series of elementary operations. The process of breaking it down required considerable mathematical insight, but thereafter carrying out the steps required only discipline.³¹ As one historian tells of Babbage:

In 1812 he was sitting in his rooms in the Analytical Society looking at a table of logarithms, which he knew to be full of mistakes, when the idea occurred to him of computing all tabular functions by machinery. The French Government had produced several tables by a new method. Three or four of their mathematicians decided how to compute the tables, half a dozen more broke down the operations into simple stages, and the work itself, which was restricted to addition and subtraction, was done by about eighty computers who knew only these two arithmetical processes. Here, for the first time, mass production was applied to arithmetic, and Babbage was seized by the idea that the labours of the unskilled computers could be taken over completely by machinery which would be quicker and far more reliable (Bowden, 1953, p. 8).

The history of computing extends back centuries, but it is dominated until the 1950s by *human* computers, those corps of labourers who carried out simple, repetitive arithmetic by the mediation of marks on paper, tools like the *abacus*, and later mechanical calculators like those of Pascal and the Hollerith (eventually IBM) Corporation. This was one of the many varieties of occluded women’s labour which sustained more widely known achievements in mathematics and science, particularly in

production model of IBM digital computer, the 701, nearly copied its design (G. Dyson, 2012). Both these machines more fully realized the stored-program concept. As Willis Ware recalls, the ENIAC could be put to use for a wide range of problems even though it was originally purpose-built for computing firing tables, but ‘reprogramming’ it still involved “physically establishing a large number of connections by means of switches, patch cords, transmission lines and so on” (Ware, 1953, p. 5). These new machines were conceived as general-purpose scientific machines from the start, and they carried out the necessary internal ‘rewiring’ with automatic switches and relays, requiring little manual intervention in the normal course of operations apart from feeding in cards with new data or instructions.

³¹ This also maps onto the syntax/semantics opposition: computing operates on symbols, but as with Shannon’s information theory, it has no concern with their semantic ‘content’ or meaning (Ekbia, 2008).

wartime as needs for calculation increased and the male labour pool sharply decreased.³² Throughout the development of the ENIAC, and later EDVAC, some two hundred young women worked on the project as human computers, with the importance and complexity of their labour as ‘ENIAC girls’ often minimized in the fawning press this new ‘electric brain’ and its male engineers were starting to receive (Light, 1999).

The ideal, as Babbage had envisioned, was to improve on the unreliable human computers through automation, but the early machines had their own myriad pathways to error, and the labour of these women along with many other engineers, operators and mechanical calculators was absolutely necessary to program the machine, keep it running, and ‘check its work.’ This was computing as ‘big science,’ in the size of its hardware and its power consumption, as well as in its mobilization of human and technical resources; it was commensurately put to use by the emblematic ‘Big’ research program, nuclear weapons research. Though the media hype which once spoke of cybernetic machines in analogies with the brain was indeed shifting toward these new machines, using them to study the brain seemed at first conceptually and materially impracticable. These were number-crunching machines, and grossly expensive ones. Yet von Neumann would direct his final writings toward exploring the relations between computer and brain (Von Neumann, 2012), and before long Newell and Simon were constructing models of mind on equally powerful, universal, but dramatically smaller and cheaper computing machines.

The fundamental unit of digital computing is the binary digit or ‘bit,’ representing the bare minimum of information, a selection from two possibilities. The material devices corresponding to this conceptual unit are switches, first the vacuum tube and then the solid-state transistor, which

³² In the First World War, for that matter, many young mathematicians including Norbert Wiener worked as human computers at the Aberdeen Proving Ground, computing artillery firing tables by hand (Edwards, 1997, p. 45). The recruitment of women for this work occurred then as well, but scaled up in the Second World War (Light, 1999).

represent bits by transitioning between ‘on’ and ‘off’ states depending on the input current. In the original, abstract ‘paper machine’ formulation of the Turing machine, there was only the processing unit and the indefinitely long paper tape. The tape was simultaneously memory, data, and program; given only instructions for manipulating the tape and the markings thereupon, and for reading new instruction sets from the tape, this hypothetical machine could carry out any computation. The paper where Turing introduced the idea sought primarily to operationalize the notion of ‘effective computability:’ it had long corresponded with the kinds of tasks which Babbage saw as readily carried out by ‘unskilled computers,’ but those computers in fact brought to bear a considerable training and mathematical intuition when applying their instructions. The Turing machine could do no such thing, and so if it was computable by such a machine, it was effectively computable, period, and *vice versa* (Turing, 1936).³³ Though Turing himself was quite interested in the construction of actual machines, his original presentation of the machine was wholly mathematical and abstract, stripping away everything but the essence of computation, as readily carried out by a human with pencil and paper working ‘mindlessly’ as by a machine—though it might take the poor human an eternity given a complex enough task.

In practice, however, the design of computing machines typically involves the interconnection of multiple different kinds of systems. The von Neumann architecture was directly inspired by McCulloch and Pitts’ work, along with other cybernetic elaborations of the neural-network concept (Conway and Siegelman, 2009; Dupuy, 2000). Developed in the course of experimentation with the Institute for Advanced Study (IAS) computing machine, it employed one

³³ In this paper Turing also answered David Hilbert’s *Entscheidungsproblem* in the negative, showing contrary to Hilbert’s own conviction that, just as Gödel proved the essential incompleteness of logic and mathematics, even with infinite time these universal computers would be unable to decide on the truth or falsity of certain propositions. With the publication of related proofs by Alonzo Church, this eventually came to be known as the Church-Turing thesis: functions are ‘computable’ if and only if they can be carried out by a Turing machine, and for any such machine there will be some *non*-computable functions.

system of electronic vacuum-tube switches and relays to constitute the central processor, and two distinct storage systems.³⁴ The first was a short-term, capacity-limited electronic memory, first planned around the RCA Selectron vacuum tube, but when supply of those tubes failed to materialize soon enough, they implemented a clever repurposing of widely-available cathode-ray tubes: the Williams tube, first developed in Manchester. These were intended as radar displays, but the scanning beam could be used to store information in ‘charge wells’ that persisted on the surface of the tube for a fraction of a second. Like modern random-access memory, they required continual power not only to modify the contents but to ‘refresh’ the existing data. The other storage medium was much more of a constant across the history of computing, from its prehistory in the Jacquard loom, through Turing and the Hollerith Corp. to roughly the 1980s: punched paper tape. This form of long-term storage required no ongoing supply of electricity and was cheap enough to be effectively infinite, just as Turing imagined, but it was also painfully slow and unreliable to process.

The material forms taken by storage technologies over the years varied substantially, from Williams tubes to acoustic delay lines, magnetic drums, and myriad others, all differing in speed and reliability, making for machines with different operating characteristics. A central aim for the new approach to modeling cognition, however, was to step away from the materialities of both electronic circuits and biological neurons. The first computing machines were already developing in this direction: while at base their functioning was binary and equivalent to a Turing machine, the electronic switching components were configured in assemblies that were intended to carry out specific elementary operations like addition and multiplication. These were caused to perform

³⁴ This machine was one of the first and most important computing projects in the United States, constructed by a team headed by von Neumann, Julian Bigelow, and Willis Ware, in the rather unusual environs of the IAS—a story told in painstaking detail by George Dyson (*G. Dyson*, 2012). Its plans were freely shared, directly informing the construction not only of JOHNNIAC, but the MANIAC at Los Alamos, the IBM 701, and many others, thereby ensuring the dominance of the von Neumann architecture for many years thereafter.

sequences of actions first by operators physically changing connections and switches (as with ENIAC), then later by encoding new instructions in binary numbers on paper tape. This was ‘machine language,’ the most direct way of interacting with a computer and at the time the only one; it was “extremely difficult to use and even more difficult to debug” (Edwards, 1997, p. 247). In 1949, the designers of ENIAC introduced Short Code, the first ‘assembly language’ that gave a more readily human-readable form to the basic instructions of the machine, but which still required that instructions be entered “in the exact form and order in which the machine would execute them,” frequently unconventional by human standards (*ibid.*). Programs were thus easier to compose, but still demanded engagement with the details of hardware implementation.

The rise of symbolic modeling in cognitive science presupposed the development of a further level of technical abstraction, though, in the form of higher-level languages like IBM’s FORTRAN, Newell and Simon’s Information Processing Language, and McCarthy’s LISP, along with the compiler programs which automate the process of translating their instructions into machine language. The field would eventually take up natural-language processing and translation as some of its central research problems, but in fact it depended upon this more fundamental feat of machine translation: constructing a compiler, the ‘great-grandfather of all programs,’ tasked with foreseeing all the operations used in any sort of computation, and their machine-code equivalents (Philip Morse, in Edwards, 1997, p. 248). This comes with clear advantages, most notably the ability to write programs which could be run on multiple different systems with appropriate compilers, as opposed to the assembly languages which were specific to each machine architecture—this ‘portability’ of software was another technical prerequisite for the conceptual development of multiple realizability. The rise of high-level languages was far from an obvious step, however, but rather the focus of a sharp controversy, with hardware experts ‘horrified’ at both the inefficiency of

running programs twice—using scarce computer time first to run the compiler and then the actual program—and the required complexity of compiler programs, which at first produced unreliable code (*ibid.*).³⁵ Eventually, as the speed of the machines rapidly increased and compilers improved, the increased computational cost of high-level languages became less of a burden. They proved valuable for expanding not only the ranks of potential programmers, but the kinds of programs which could be envisioned.

Computer science as a whole was thereby shifting to a paradigm of general-purpose symbolic representations, rather than narrowly mathematical ones. For the purposes of cognitive science, one task domain had immediate appeal: formal logic. Proving theorems in logic was emblematic of the highest achievements of human cognition, and since the time of Aristotle had been understood as a distinctively human mark of intelligence. Mathematicians like Bertrand Russell, Alfred North Whitehead, and Kurt Gödel had laboured enormously over the past decades to demonstrate the foundations of all mathematics in newly developed formal systems of propositional logic, and were rightly regarded as geniuses for having succeeded. The work of Turing and Church showed that any mathematical or logical operation for which we could specify a finite set of instructions – an algorithm – could be carried out by a simple machine with only a few elementary operations. Turning complex mathematics and logic into algorithms, however, at first demanded enormous ingenuity and hands-on experience with the hardware of computing machines. Indeed, in its infancy cognitivism was nothing more than the set of psychological researchers with access to and experience working with digital computers. Considerable portions of this labour were proving susceptible to automation, however, first with assembly languages, then compilers and higher-level

³⁵ Along similar lines, Newell and Shaw noted that there were a range of shortcomings that revealed themselves after about fifty hours of experimentation with their first iteration of the Information Processing Language, foremost among these the cost in memory space and time of the language and its compiler (Newell & Shaw, 1957, p. 232)

languages. Herbert Simon and Alan Newell thus saw it as an evident next step to try writing a program which might automatically carry out the process of actually deriving and proving theorems of logic, mechanizing what Russell and Whitehead had done with great mental effort aided only by pencil and paper just a few decades prior.³⁶

Given the fundamental analogy recognized since even before cybernetics, linking electronic switching systems and systems of neurons, along with the relatively novel understanding of mathematics as grounded in formal logic, a synthesis of all these ideas in software seemed eminently justifiable.³⁷ Brains could be seen as ‘functionally equivalent’ to computers without any notion of deep structural similarity, or need to mirror the structure of the brain in the design of machines. Newell and Simon set out to construct “a theory of the information processes involved in problem-solving and not a theory of neural or electronic mechanisms for information processing,” leaving the latter as a task for “another level of theory construction” (Newell, Shaw, & Simon, 1958a). A great many independent projects were undertaken along these lines in subsequent years. In the remainder of this chapter, I consider a few important examples of these tools for thought, and equally the programming tools which they relied upon, starting with Newell and Simon’s Logic Theorist, Logic Language, and Information Processing Language. Even based on the earliest ‘hand simulations’ of this program (which I describe below), they were firm believers in its capacities, contending that it could prove “most of the 60 odd theorems in Chapter 2 of *Principia Mathematica*” (Newell & Simon, 1956, p. 79). Simon reportedly told a graduate class at the beginning of 1957 that “over Christmas, Al Newell and I invented a thinking machine” (Crowther-Heyck, 2005, p. 1): they’d written a

³⁶ Before Logic Theorist, they had already been at work on another paradigmatic AI task, the playing of chess, with Newell’s inspiration coming from the prior work of Oliver Selfridge. Though the symbols and rules for manipulating them are quite distinct, chess and logic are both readily implemented in software, as discussed below.

³⁷ This is a direct sense in which cognitive science participates in the broader ‘organizational revolution’ (Heyck, 2014), as discussed in the previous chapter.

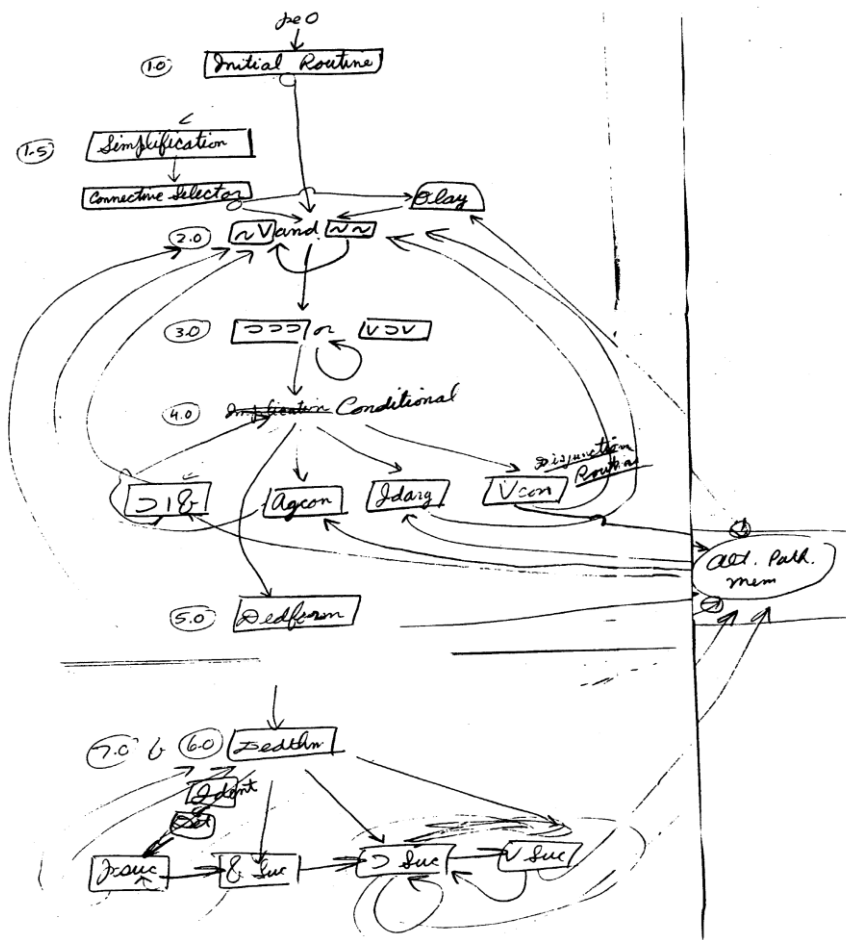
computer program “capable of thinking non-numerically, and thereby solved the venerable mind-body problem, explaining how a system composed of matter can have the properties of mind” (Simon, 1996a, p. 190). Bertrand Russell was reportedly ‘delighted’ to hear of Logic Theorist’s successes, jokingly lamenting that he wished “Whitehead and I had known of this possibility before we both wasted ten years doing it by hand” (Simon, 1996a, p. 208).³⁸

The task for Logic Theorist (LT) was straightforward: to “prove that certain expressions are theorems—that is, that they can be derived by application of specified rules of inference from a set of primitive sentences or *axioms*” (Newell & Simon, 1956, p. 67). Though digital computers are in many ways perfectly suited to carrying out logical operations, deriving theorems is not a straightforwardly computable problem in the same way as calculating logarithms or firing tables. First of all, in what its creators saw as its greatest innovation towards ‘thinking,’ it was “truly non-numerical,” operating instead on propositions in the formal language of Russell and Whitehead (Newell & Shaw, 1957, p. 231). The Logic Theorist program included two basic connectives, a third derived one—conjunction, disjunction, and implication—five axioms, and two rules of derivation, all derived from *Principia Mathematica*. It would be given a more complex expression, for which it would then attempt to construct a valid sequence of steps demonstrating, if possible, its proof from these first principles. In theory, it is possible to carry this out with a ‘brute-force’ or exhaustive search algorithm: the program could simply try every possible permutation of its basic axioms given its rules of derivation, until it ended up with the ‘goal’ expression. This would be grossly inefficient,

³⁸ In an exchange which Simon reproduces in full in his autobiography, Russell goes on to state that he is “quite willing to believe that everything in deductive logic can be done by a machine,” and expresses pleasure that the LT devised a proof distinct from and in some ways more elegant than the *PM* version (Simon, 1996a, pp. 207–210). The whole correspondence is quite lighthearted, and I say Russell was evidently joking above despite his clear beliefs about the logical capacities of machines, since both well understood that whatever its successes in proving theorems, the task of independently devising an entire system of logic was well beyond the capacities of LT.

however, simultaneously infeasible given the computing hardware of the day, and not a particularly good model of human problem solving.

Instead, what was required were ‘rules of thumb,’ or what Simon and Newell would call *heuristics*: different sorts of algorithms which didn’t specify the entire process in detail, but were useful guides toward solutions. Partly based on his empirical investigations of human problem solving, and partly on his own introspection with respect to theorem-proving, these rules were intended to reproduce some of the higher-level thought processes that human minds bring to bear



in these same tasks.³⁹

Mathematicians, even early students of logic, don’t exhaustively try every possible permutation of their formulae, but intuit what kinds of steps are needed as they develop a feel for the medium. LT tries to formalize and automate the cues upon which this logico-mathematical intuition

Figure 15 - Logic Theorist sketch

³⁹ These empirical investigations included his earlier experience at the RAND Systems Research Laboratory, examining the workings of a model Air Force SAGE station—which I discuss further below—and more directly in university psychological laboratories, where they asked college sophomores to solve symbolic logic problems, “thinking aloud” as they worked (Newell, Shaw, & Simon, 1958b, p. 2).

operates. If a proof went on even five or ten steps, then exploring all possible transformations of the propositions by trial-and-error search would demand thousands of branching paths, each representing potential ‘proof chains.’ But if LT could respond to cues that allowed some of the options to be eliminated from consideration at each step, as likely dead ends, then that number could be dramatically pared down to a few dozens. Like the machines of cybernetics and ‘classical servo theory,’ LT functioned by a feedback mechanism (Figure 15), but instead of replicating the signals flowing through the nervous system, it sought to replicate the ideas flowing through the mind of a human problem solver, with distinct ‘housekeeping,’ ‘substitution,’ and ‘matching’ subprograms, among others, as well as working and storage memory and an ‘executive’ routine. Equally it abstracted away from the circuitry of particular computing machines, requiring the creation of a higher-level language, which they called the Information Processing Language, and presented in simplified form as ‘Logic Language’ (Newell and Simon, 1956). While Newell, Simon, and their collaborator Clifford Shaw all regarded the others as equal partners in the creation of LT, Simon was the primary designer of the program itself, while Newell and Shaw were more responsible for the creation of the language and compiler.

Simon and the new cognitivist group saw this simulation of information-processing psychology—Logic Theorist, and programs like it—as far more plausibly claiming the title of ‘thinking machine’ than any of the cybernetic-era machines. This was paradigmatic of the tools-to-theories heuristic which I seek to highlight here, inasmuch as Simon explicitly stated that he was led to formulate this conjoined psychological model, theory, and AI system by his experience working with computers at the RAND Corporation Systems Research Laboratory (Simon, 1996a). It was also the starting point for heuristics themselves as the core of the cognitivist approach, as its primary means for ‘pruning’ searches and using algorithms to tackle problems which do not come well-

tailored for computing. The first electronic digital computer to execute LT was the RAND Corporation JOHNNIAC, a successor machine to the IAS computer, whose construction was supervised by von Neumann at the Air Force-funded think tank, where Newell and Simon had met and were employed at the time. Yet the program was architecture-independent indeed, and its first functional realization was a family affair, “nature imitating art imitating nature,” as Simon assembled Newell, his wife, three children, and several graduate students together in a campus building on a winter night, with cue cards representing English translations of the operating rules for LT subroutines. Each member of the group followed the instructions on one card, storing logical representations with pencil and paper on different cue cards designated as ‘memory.’ Simon recalls in his autobiography that “the actors were no more responsible for what they were doing than the slave boy in Plato’s *Meno*, but they were successful in proving the theorems given them” by rigidly following the prescribed rules (Simon, 1996a, p. 207). It was in fact on the basis of this manual, human simulation of the Logic Theorist’s simulated reasoning that Simon made the initial proclamation of having created a thinking machine with AI Newell over the Christmas break in 1955: they’d literally created it by hand with paper and pencil, ‘programming’ family members and students to follow its instructions.

It was not until August of 1956 that the program was actually implemented on JOHNNIAC, with its full specifications published later in the year. Ultimately the program functioned much like a novice student of logic. Rather than having a clear intuition of how to prove theorems, it depended considerably on rote application of rules by trial-and-error. The program’s executive routine ran through a list of methods by which the basic axioms could be transformed (adding in additional variables or employing valid substitutions), deciding which of those produced an expression that was most similar to the desired theorem, and trying some modifications on that, periodically checking its

results against the target in a feedback loop and storing intermediate states in memory. If the expression continued to transform into something more like the goal theorem, LT would keep following that branch of the search tree, running through all its potential operations until a proof was found; if a ‘dead end’ was reached, it would step back to its memory of a previous state, and try another permutation that wasn’t ranked as highly by its initial rule of thumb (Gardner, 2008).

Though it was proposed primarily as a ‘performance program’ rather than a ‘learning program,’ and its actual program does not change in the course of operation, learning does take place in one “very important” respect, in that it stores theorems as it proves them, and then calls upon these as “building blocks” for the proofs of subsequent theorems (Newell & Simon, 1956, p. 74). The program was specifiable only with the innovation of Newell and Shaw’s list programming language, the first of its kind, which in turn allowed for the implementation of its hierarchical structure, relying heavily on “both iteration and recursion” (*ibid.*, p.78).

Rather than requiring as many subroutines and program elements as there were potential components of a logical expression, it decomposed the expressions into hierarchical tree structures and individual elements—corresponding in implementation with a single 40-bit ‘word’ of

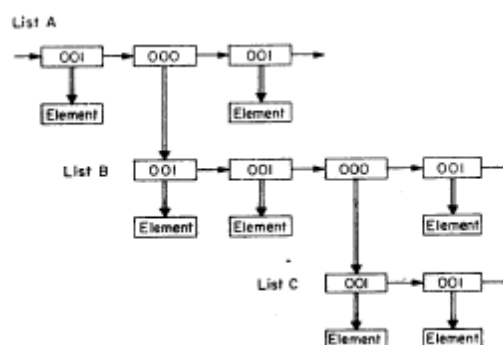


Figure 16 - List programming

JOHNNIAC memory—and it was able to call these routines for matching, substitution or decomposition on elements at any level, then recursively once again on their outputs if necessary (Newell & Shaw, 1957). ‘List processing’ meant that each element would contain the ‘address’ of the next, and that elements could themselves point to other lists, all of which could be dynamically remade. Elements in memory could thus be stored and processed in a variety of ordered list

structures, not directly tied to their physical locations in hardware. Without these programming techniques, Logic Theorist (along with most other modern applications of computing) would require impossibly, infinitely many lines of code.⁴⁰ Heuristics, recursion, and higher-level list programming were the core technical underpinnings for the theoretical innovations of symbolic cognitive science, all present within LT. Newell and Simon (absent Shaw) first presented their model at the 1956 Dartmouth workshop, and quickly thereafter disseminated its design and functional properties to the whole circle of researchers involved. Claude Shannon and Oliver Selfridge, Wiener's former assistant, were in attendance, along with Marvin Minsky and John McCarthy. Shortly thereafter in September, a 'Special Interest Group in Information Theory' was also convened at Dartmouth, and Newell and Simon presented their program once again to a group which included many of the same actors, along with George Miller and Noam Chomsky (Miller, 2003). This completed the core set of cognitivist researchers, and enshrined 1956 as the year of its proper origin. To Newell and Simon's surprise, however, the reception of their program at the Dartmouth events was a bit frosty. They themselves recall that there was likely an element of arrogance in their view that they were the sole researchers in attendance with a truly concrete and functioning example of a 'thinking machine' (Newell, 1991; Simon, 1996a).

Simon quickly established himself as the most ambitious proponent of this new vision for psychology and artificial intelligence as a joint enterprise. He believed that Logic Theorist was

⁴⁰ The reproduced discussion of Newell and Shaw's paper shows that recursion and the potential infinite regresses or paradoxes to which it gives rise were a particular interest of those in attendance: the first question asks "do you have a list of those lists which do not list themselves?" garnering the reply from Shaw that "there is not a direct answer to this question, but the debugging list does list itself." Another asks for an example of recursion, a 'subroutine using itself,' and Shaw notes that the matching routine does so, recursively looking at main connectives in the starting expression and seeing if they match the desired theorem, then looking at each sub-element and carrying out the same operation. Upon terminating one of these trees "it says, 'I am done' to itself, reiteratively, and then backs up to a certain point," at which it proceeds down the other 'right' branch; if everything matches, then it has successfully proven the theorem, if a mismatch is found, it calls one of its routines for modifying that sub-expression (Newell & Shaw, 1957, p. 240).

literally a ‘thinking machine,’ based on his concept of thought as developed with Newell: the physical symbol systems hypothesis. This was the view that any system capable of adaptively performing operations and transformations of symbols—that is to say representations of the world, though as for McLuhan and information theory, their content was less important than their form—was a system capable of *thinking*. Thought simply was symbolic computation, whatever the material characteristics of the system in which it was instantiated. Though they of course recognized that the capacities of humans were far beyond those of LT, Newell and Simon believed that by programming a system to carry out ‘non-numerical,’ general-purpose symbolic computation, they had created something that was now only quantitatively rather than qualitatively different from the human mind. All that remained was to develop list-structured databases and heuristics appropriate to all the fields of human endeavour, and the computer hardware capable of storing and processing them. No small task, but Newell and Simon were (in)famously confident in the progress of their field, arguing that on the basis of their early successes in heuristic programming, using processes ‘closely parallel’ to human problem-solving, they could predict:

1. That within ten years a digital computer will be the world’s chess champion, unless the rules bar it from competition.
2. That within ten years a digital computer will discover and prove an important new mathematical theorem.
3. That within ten years a digital computer will write music that will be accepted by critics as possessing considerable aesthetic value.
4. That within ten years most theories in psychology will take the form of computer programs, or of qualitative statements about the characteristics of computer programs. (Simon & Newell, 1958, p. 7)

While these were the most dramatic and specific predictions made, the entire cognitivist group saw the approach of heuristic programming using higher-level languages as bound to yield increasingly intelligent programs and significant insights into human cognition.

At the same time, John McCarthy was working on techniques for time-sharing with computers. One early computer, for instance, was at first shared between IBM, MIT, and BBN,

taking one 12- and two 7-hour shifts, respectively. Rather than having each research group monopolize a computer for a stretch of time, with this new software processing power would be dynamically and continuously shared. This would allow each group to input their data and check their outputs separately, removing the major bottleneck, and the computer would allocate its processing time according to set priority rules as different instructions came in. This innovation was also tied to networking, in that remote terminals in different facilities were a useful complement. At the time, there wasn't much to share in terms of computing power, but McCarthy and eventually Licklider saw that time-sharing would become increasingly important as computers' power and applications grew (while they remained few in number). There was no necessary or direct connection between AI and time-sharing, but it did permit greater access to "the essential tools of their trade" for AI researchers, and more importantly McCarthy saw it as providing "*the right subjective environment*" for their work (Edwards, 1997, p. 258): it created a world in which it made more sense to speak of a dynamic, *interactive* relationship between humans and computers, feeding in instructions step by step and seeing their outputs, rather than constructing long strings of punched cards representing machine code and waiting for output in a similar form.⁴¹ When McCarthy saw Newell and Simon's presentation on LT and its programming language, he also immediately recognized how well it fit with his other two interests, in time-sharing and AI. Though he recounts there was 'little temptation' to copy their IPL, based as it was on a "JOHNNIAC loader that happened to be available to them," McCarthy soon set about constructing his own list-processing language on the 'algebraic' model of

⁴¹ Joshua Lederberg, among the creators of the later DENDRAL heuristic program, describes how the 'serial batch mode' of earlier computers produced a subjective and collective environment with its own pros and cons, noting that he would often need to work late at night to ensure access to computing time, and that "the democracy and night-owl ambience of the batch system was a social mixer for several enthusiasts from wide-ranging disciplines." Only the most dedicated were gathered at the midnight shift, and workers in that era necessarily spent a great deal of time trying to "simulate the machine" in their own thought, as opposed to the "casual, experimental mode—'let's see if this works'—of today's interactive systems" (Lederberg, 1987, p. 5).

IBM's FORTRAN (McCarthy, 1996).⁴² All of these techniques were appealing to ARPA, which would go on to be the major funder of cognitivist AI research throughout this period. Joseph Licklider argues quite explicitly that the interests of science, military and industry were inseparable in this period: that “what the military needs is what the businessman needs is what the scientist needs,” and that you can make the exact same project appealing to all three sets of actors provided it has several ‘facets’ (Licklider, 1988, p. 39). Yet other scientists, like Wiener, were quite openly questioning this alignment of interests in the Cold War era. Rather than describing a ‘natural’ coordination of goals, the phenomenon Licklider describes should instead be viewed as his signature accomplishment: the engineering of these heterogeneous interests into a unified funding program.

While cognitivism developed into the core of a new mainstream in psychology, it retained a certain essential heterogeneity and vagueness, held together still more by tools than theories. It flourished because of, rather than in spite of this pluralistic character: as Licklider put it, “it was not a clear vision, as certainly, not that ‘We’ll plug them all together, and that will be the system;’ but, rather, it was a matter of getting a supply of parts and methods and techniques, and different people will put together different systems out of it” (Licklider, 1988, p. 43). Licklider took the position of Director at ARPA’s Information Processing Techniques Office in 1962, a position he held for just two years. Over the course of the 1960s, however, the agency funded the development of several

⁴² In a fascinating (though peripheral, for my purposes) turn of events within the history of AI, IBM was one of the major early supporters of AI research in this period, funding some of the first serious efforts at computer chess-playing. In 1959, however, after employing McCarthy as a consultant on a list-processing extension to FORTRAN for three years, there was a ‘purge’ of all things AI-related at IBM. McCarthy recalls that “there was allegedly some PR aspect to turning it off. Namely, IBM thought that artificial intelligence was bad for IBM’s image -- that machines that were as smart as people and so forth were bad for their image. This may have been associated with one of their other image slogans, which was ‘data processing, not computing.’ That is, they were trying to get computing into business, so they wanted it to look as familiar and unthreatening as possible... IBM didn’t officially do any AI between 1959 and 1983. That’s when they revived it” (McCarthy, 1989, p. 10). That late revival led to the vindication (far behind schedule) of Newell and Simon’s first prediction, with the 1997 victory of the Deep Blue supercomputer and its chess-playing program in a match with Garry Kasparov. I return to IBM’s contemporary work in AI and ‘cognitive computing’ in the following chapters.

graduate programs in computer science (as discussed above), along with a range of basic-research projects in artificial intelligence and interactive computing. With RAND Corporation, the IPTO funded the boom in cognitivism. Simon and Newell were making dramatic predictions, and Licklider was a ‘true believer in his own propaganda’ (*ibid.*). McCarthy and Minsky at MIT were given a large grant to establish ‘Project MAC,’ an experimental time-sharing computer with many distributed consoles around Cambridge, and a realization of some proposals from ‘Man-Computer Symbiosis.’

Newell and Simon also expanded on their Logic Theorist concept with the ‘General Problem Solver’ program. This was specified in the fifth iteration of Newell and Simon’s IPL, and now logic was merely one of the ‘task modules’ (along with the “Missionaries and Cannibals” river-crossing puzzle) for a general-purpose heuristic program, which could be programmed with new tasks by

Goal type #1: Transform object a into object b

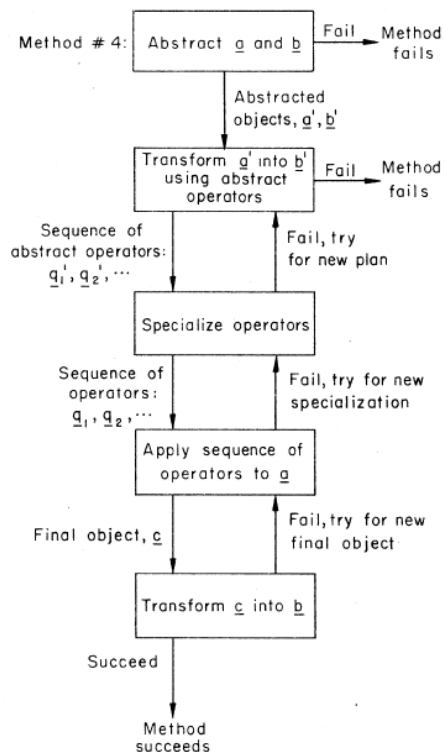


Figure 17 - Planning diagram

decomposing them into recursive applications of two basic techniques, means-ends analysis and ‘planning.’ Planning, here, meant the application of an iterative feedback process like that depicted in this diagram (Figure 17), and the application of heuristics to determine which components of the symbolic expressions stored in memory should be operated upon (Newell, Shaw, & Simon, 1958b, p. 21). GPS dealt only with the ‘most general features’ of a given task environment, and decomposing tasks into these general features in such a way that the program could accept them

required considerable human labour (Newell, 1963).⁴³ The software's planning process constitutes a step of further abstraction, where it attempts to simplify the problem (for instance by ignoring different logical connectives to look only at arrangements of propositions), recursively applies its means-ends analyses to solve the problem in this simplified form, transforming the input state 'a' into the goal state 'b,' and then applies this same sequence of operations on the real task. All of these elements of GPS were intended primarily as models of human psychology by Newell and Simon, though they were equally convinced of their possibilities in practical applications. For the time being, however, GPS and other 'thinking' programs remained limited to 'toy' problems.

Among these was chess-playing, which eventually came to be known as the *Drosophila* of AI for its central role as a model—in this case a model *problem*, in lieu of an organism. Claude Shannon and Alan Turing had already published some early forays into programming computers to play chess. GPS could be set to the task as well, though it did not come pre-programmed with such a module, and it never saw great success. Newell and Simon observed, though, that it could readily be put in a form amenable to GPS, of changing a state 'a,' using the rules of transformation specified by the legal moves in chess, to a state 'b' where one's pieces have the other side in a state of checkmate. Transforming the problem into a symbolic representation is relatively easy—human chess players themselves use a notation for representing piece positions and moves with short, unambiguous strings of numbers and letters. As in formal logic, in principle the problem was amenable to pure brute-force search, and many current programs employ this method, given their access to exponentially greater computing power. There are many rules one could apply as a rough guide to

⁴³ There is evidence for this in the fact that GPS in its published version of 1958 included only the original logic task and the 'Missionaries and Cannibals' puzzle, though later puzzles like the Tower of Hanoi were made into forms solvable by GPS. Harry Collins has made the case that this extensive human labour invalidates the notion that these programs were truly 'thinking' in an autonomous and intelligent way (Collins, 1990), an argument that certainly holds true against the strongest claims of Newell and Simon for their programs, but less clearly so against the possibility of artificial intelligence in general.

prioritize the search tree, however, excluding certain unlikely branches. Developing these heuristics was the difficult task in programming a computer to compete successfully against a human being, but it was amenable to a variety of approaches borrowing from empirical investigations and introspection by the programmers (who nearly all played the game themselves with some skill). The problem had obvious appeal, and in many ways has long served as the most public face of artificial intelligence research. The core of cognitive science research has been in decidedly less media-friendly programs, however, particularly as the computing power became available to tackle chess with simpler, non-heuristic methods. The IPL was eventually outpaced by McCarthy's Lisp as the main programming language for AI and cognitive modeling. Lisp functioned on the most widely available machines of the time, the Digital Equipment PDP series and the IBM 700-series, and was intended to facilitate coding programs in terms of goal/feedback structures, plans, and dynamic lists. Since all these lists required storage of the address for the subsequent list item in addition to the content of each element, they were highly inefficient on standard computers, and so Lisp in turn drove the development of specialized 'Lisp machines' optimized to compile and run the language.⁴⁴

Another notable project of the MIT AI Laboratory, which developed out of the ARPA-funded Project MAC, was SHRDLU, a program by Terry Winograd written for a PDP-6 computer, which attempted to apply these same heuristic programming techniques to the problems of both natural-language processing and purposeful action in the world. Its 'block world' became one of the most widely known applications of Lisp and symbolic AI. SHRDLU would accept instructions and questions in English about blocks, and in response it would output English sentences, performing

⁴⁴ Though the language remains in wide use and several manufacturers sold Lisp machines in the 1980s, this attempt at commercialization (by an offshoot of Project MAC and the MIT AI lab) was ultimately unsuccessful, as computers able to run the language just as fast without specialized hardware started to enter the market.


The Robot's World

Figure 18 - SHRDLU's World

requested actions within a computer-generated virtual space with blocks of various shapes (Winograd, 1971, p. 25). Thus one could ask it to 'pick up the big red block' and it would do so; if there were multiple possibilities it would ask to clarify which block one meant; users could define certain combinations of blocks with novel terms, and it would remember them; it would manipulate the blocks as requested, and an independent sub-program would simulate their actions.

A user could even tell it 'I like the box,' for instance, or 'the blue pyramid is nice,' and it would remember that data, returning it later if asked directly ('is the box nice?'; Winograd, 1971, p. 41). This was a highly successful demonstration of AI for its seeming ability to link together language processing, and interaction with both user and a 'world.' As with most other projects that came to be seen as emblematic of cognitive science, Winograd's efforts focused on the translation between mind and machine rather than either pole exclusively. In the tradition of the AI lab and Project MAC, he likewise wove together several distinct interests: in the understanding of mind, in the creation of autonomously intelligent systems, and in the development of new techniques to facilitate human-computer interaction.⁴⁵ But despite the high hopes many shared with Newell and Simon, when these techniques were applied to problems outside the simplified domains of chess-playing or

⁴⁵ Licklider recounts that he helped push for all three of these from his own interests, and because he was simply one of the few people spending several hours at a computer console daily. Thus he was intimately acquainted with both the possibilities and limitations of computing at the time, though he did not acquire a console in his office until near the end of his tenure at ARPA. He recounts another driver of interest in improving interactive computing: they found that consoles were best placed in small rooms without space for more than one person, "because if an admiral were sitting at the console, and there were junior officers looking, the admiral was afraid to move his fingers for fear he'd reveal he didn't know what to do" (Licklider, 1988, p. 56). For more on Winograd's later career and his take on Licklider's 'IPTO hype,' see *n.*62 below.

virtual worlds, to the kinds of messy, noisy situations where military funding agencies hoped to apply them, there was no certainty of success.

Cognitivism was a movement which stepped away from the materialities of cognition, drawing inspiration from psychology and digital computing to see itself as decisively overcoming the *metaphor* of mind-as-computation by rendering it entirely literal. Computers still functioned as models and metaphors, of course, but for those who accepted cognitivism in Newell and Simon's formulation, computers were also instantiations of the same fundamental process which produced human cognition. Minds and computers were understood on equal terms as physical symbol systems: implemented in different ways, the former evolved and the latter designed, but capable of converging functionally without mimicking the biological details. Systems such as formal logic, human reason, and Turing-equivalent computers, which modified representational symbols according to algorithms and heuristics, all 'exercised intelligence' by searching through possibilities for changing one set of symbols into another target configuration—and they regarded this as fully demonstrated, speculating that the most interesting further question was whether there were significant commonalities between information processing mechanisms in minds, computers, and the genetic code (Newell & Simon, 1976). I discuss the failures of the cognitivist programme in the next section. Simon, for his part, was not so much 'wrong' about having created a thinking machine, nor did he simply carry a metaphor too far. Rather, he sought to redefine the literal sense of cognition in a way that encompassed certain behaviours of digital computers: any system capable of processing symbols, rather than merely numbers, using heuristics, rather than 'pure' algorithms, was thus a thinking system. Where any number of other scientific fields seek to construct models, cognitivism elevates them to an equal plane with the phenomena themselves, and makes symbolic media its primary explanatory principle. Given that none of his predictions were successful within

the timeframe he proposed, we may regard his programme too as falsified in some sense. Yet as long as it seemed to be delivering results within its stated task domains, and as long as there were sympathetic actors like Licklider controlling funding sources, Simon's physical symbol system hypothesis—his new definition of cognition—seemed to hold. If symbols and search had delivered machines capable of matching human capacities for understanding language, perceiving the world, and behaving adaptively, no one would question the literalness of this homology between mind and computation. Efforts seemed to flag, as did the fortunes of the hypothesis, but it retains many advocates and may still deliver on its promises.⁴⁶

The concept of cognition at its most abstract and general emerged from a very specific social milieu, once again in the heart of the military-industrial apparatus. It is integral to this story, not merely a matter of funding 'spin,' that Simon's career was oriented toward economics, operational research, and military applications. Other researchers took their modelling efforts in different directions, but cognitivism was united by its digital tools, along with its closely linked belief in the centrality of symbolic representations and heuristically-guided search. I have tried to emphasize throughout that as important as metaphor may be, it is only part of the story. Metaphors are the most mobile and suggestive, but also among the vaguest elements in a reflexive feedback system, where they guide the construction of theories and tools alike; the more important thing, from the standpoint of the scientists, is the way that turning these metaphorical ideas into working programs constitutes a tangible output, which then feeds back into technoscientific practice, reshaping thinking about thinking. This process took place in many sites, but one of its most important points of departure was the RAND Systems Research Laboratory, which studied the performance of the

⁴⁶ As mentioned above, IBM's Deep Blue did ultimately defeat the world's chess champion, but it did so more by brute-force search than advanced heuristics (or any trace of reflective consciousness), and so complicates the notion that it was 'thinking' in a human sense.

Air Force's SAGE control centers (Edwards, 1997; Heyck, 2008b). It was there that Simon met Newell, and every summer throughout the 1950s he returned to work there on simulating the SAGE system and understanding the many complex interactions between humans and computers it required. The SRL was their first time extensively working with digital computers, and their first opportunity to observe a whole human-machine assemblage in detail. Its creators, Robert Chapman and John L. Kennedy, saw the entire laboratory as an integrated, highly complex 'organism,' whose components could be understood in separation, but whose overall function demanded a new understanding. By considering a model made of "metal, flesh, and blood," of many humans and machines rather than one individual, the process is 'magnified,' and the communications channels that are hidden within the typical organism are 'brought into the open' (in Heyck, 2008a, p. 55).

Work at the Laboratory indicated that such assemblages—'cyborgs' *avant la lettre*—could indeed learn and function as an integrated unit, but that there were considerable challenges in making them function as expected. Central to their collective operation was the development of heuristics, allowing the organization to "spend its efforts more effectively by determining what efforts will be given priority," picking out and focusing on the most 'potentially threatening' radar tracks rather than all of them (Kennedy & Chapman, 1955, p. 9). Though Simon and Newell went on to model the individual 'thinking' unit in their work with LT and GPS, their framing of cognitivism in no way implies a world of atomistic information-processors. Rather, Simon's conception of the mind is entangled from the start with social institutions. These institutions, as distributed and collective information-processors—be they military bureaucracies, universities, or corporations—form the 'givens' which make effective reasoning possible given our limited ('bounded') powers in isolation. The "rational individual," for Simon, "is, and must be, an organized and institutionalized individual" (Simon, 1997, p. 102). Thus he can be read as presenting something

of a Foucauldian argument, but with radically different conclusions: from a general analysis of the principles by which institutions shape and govern our thought, Simon hoped to understand how to improve upon them, pursuing maximally rational decisions, efficient distributions of resources, and stable social order. Both regard power as ‘productive’ (Foucault, 1995). The effects of our institutions are not mere constraints or prohibitions, an ‘iron cage’ of bureaucracy (Crowther-Heyck, 2005, p. 117), but rather they are the very conditions of possibility for our rationality.

Cognitive science as such was finally formed through the exchanges between Simon and Newell’s group at the SRL (then later Carnegie Mellon), the AI Lab at MIT, and the more traditional psychological research of the Harvard Centre for Cognitive Studies. The work of this latter group is more typically recounted as constituting the ‘cognitive revolution,’ or the return of mind to scientific respectability after the benighted era of behaviorism (Gardner, 2008; Miller, 2003). George Miller first published his famous paper on the ‘Magical Number Seven,’ describing experimental investigations of human short-term memory capacity in terms of information-processing (Miller, 1956), and then with Pribram and Galanter described cognition more generally in terms of a ‘Test-Operate-Test-Exit’ feedback process analogous to Simon and Newell’s computer programs (Miller et al., 1960). While reflexivity has a complicated history in psychology (Morawski, 2005), it is crucial that the new theory of mind in cognitive science is doubly reflexive, and far from neutral and ‘objective.’ In the first place, it is an image of mind which ‘makes normal what had been normative:’ “borrowing ideas of democratic thinking from political culture and conceptions of good thinking from philosophy of science to describe humans as active, creatively thinking beings” (Cohen-Cole, 2005, p. 108).

Implicitly and explicitly, the cognition of scientists—particularly the distributed machine-aided cognition going on in their laboratories—structured their understanding of cognition in general.

Later developments in the field gave this more openly political stakes, mapping the ‘open-mindedness’ of American cognitivist psychology against perceived Soviet closure (Cohen-Cole, 2005, 2014).⁴⁷ But moreover, as I have tried to emphasize, reflexivity was both produced by and built into the computational tools of cognitive science. This was a reflexivity borne of the coupling between “scientists’ intuitions” and the “class of systems” they adopted as tools for research (Newell, 1982). Introspection and laboratory study at sites like the SRL and Psycho-Acoustic Laboratory informed the development of heuristics, programs, and new tools for programming; their functioning or failure to function fed back into the formulation of new ideas and new machine designs; periodic attempts were made at systematizing these insights into theories of mind; all this ultimately fed into new ideas about how technoscience should be funded and organized. It went by many names, and didn’t finally achieve closure around the title of cognitive science until the 1970s, when the Sloan Foundation ‘came calling’ (Miller, 2003). This led to the 1978 report cited at the outset of this chapter, and its performative depiction of the field as an interdisciplinary constellation dedicated to the study of mind. The date of conception for cognitive science was in 1956, but its christening didn’t come for two decades, by which point part of its appeal was that early predictions for ‘thinking machines’ had already failed, and the heyday of symbolic AI was over. If we couldn’t construct a machine able to think for itself, at least we could draw some conclusions about how thinking works (and what doesn’t work for replicating it). Cognitive science itself is a product of bridge-building and heterogeneous engineering, oriented more toward the design of working models in technological media, than the priorities of any conventional discipline for the study of this

⁴⁷ The public face of cognitive science as the discipline of the ‘Open Mind,’ in such works as Allport’s *The Authoritarian Personality*, persuasively argued by Cohen-Cole, presents a fascinating contrast (whose exploration is beyond the scope of this project) with Paul Edwards’ characterization of ‘closed-world discourse’ as dominant *within* the interlinked military-academic research apparatus of the Cold War era (Edwards, 1997).

complex phenomenon called ‘mind.’ For much of its history these models were more akin to those of Newell and Simon than the cyberneticians: based on symbolic representations and digital computers. Next, however, I discuss some of the setbacks of this specific programme in cognitive modeling, which led in turn to an ‘AI winter’ and overall funding decline for cognitive science, before turning in the final chapter of this section to more recent developments.

Cognitivism becomes ‘old-fashioned.’

In Perceptrons, we saw a seeming ‘falsification’ of the cybernetic programme for modelling cognition, to which an official history could point in validation of the shift to symbolic approaches. Cognitivism would never entirely go out of fashion. Rather, it experienced a slow decline following failures to deliver on heady promises of the early days. Winograd’s SHRDLU was widely hailed as a great achievement in this traditional paradigm, and in some respects the field seemed to be booming throughout the 1970s. But Licklider was no longer at ARPA, which in 1972 was renamed the *Defense* Advanced Research Projects Agency, in one of many efforts to emphasize the military applications of its programs. Once, ARPA had been appreciated for its ability to carry out impressive large projects on relatively small budgets, unlike the in-house R&D divisions of the various military branches in the US, which ‘knew how to spend money’ (Licklider, 1988). By the early 1970s, however, spending for such basic, open-ended AI research had started to decline, and there was diminishing interest in supposedly ‘general-purpose’ AI projects whose functioning was confined to ‘toy’ domains. Efforts in robotics, like Stanford University’s ‘Shakey’ robot, and a multi-sited five-year research programme (started in 1971) to develop effective speech understanding software, hoped to achieve something closer to applied success, while also shedding light on human cognitive processes. These achieved many of the goals they set for themselves, but were far from real-world

usefulness. Shakey could not reliably push one block onto another when directed.⁴⁸ The language understanding project, pursuing a longstanding dream for human-computer interaction, achieved some notable early successes in using straightforward pattern matching for words. But it soon encountered serious difficulties in separating out individual words from sentences, and found that “many words appearing in sentence contexts varied dramatically in acoustic characteristics depending on the surrounding phonetic environment” (Klatt, 1977, p. 1346). Hence directly matching spoken words with pre-recorded ‘definitions’ would not be fruitful, and instead the systems produced phonetic transcriptions, which were then mapped to sentences using symbolic networks quite different from those of the Perceptron era.

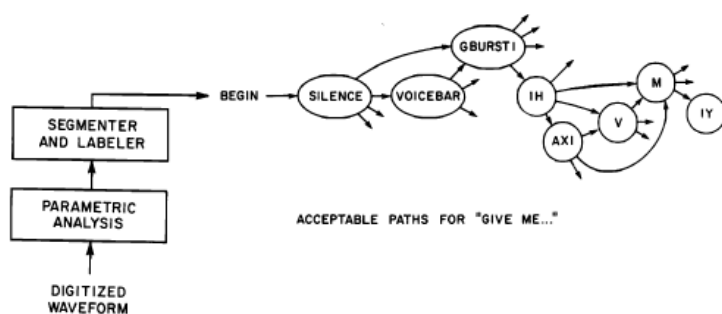


Figure 19 - HARPY Block diagram

These were ‘finite-state Markov models,’ where the nodes and arrows between them represented possible linkages between phonemes which would constitute valid English sentences.

Speech-recognition programs converted audio data into digital representations, then phonemic segments, which were then fed through a system of many such interconnected networks (as depicted in Figure 19). The most successful of these programs, the Carnegie Mellon HARPY project, was able to recognize sentences spoken by multiple speakers using distinct pronunciations, with a 95% accuracy rate. Others fared far worse. HARPY achieved its success, however, by sharply limiting its domain: it used a vocabulary of roughly 1,000 words, accepting sentences with a maximum of 8

⁴⁸ Though its model task was even simpler than SHRDLU’s virtual-world example, Hans Moravec, a graduate student at the time, recalls that “it did it in several independent attempts, which each had a high probability of failure. You were able to put together a movie that had all the pieces in it, but it was really flaky” (in Crevier, 1993, p. 115).

words, using recordings “made in a quiet room using a good-quality microphone,” and a ‘low branching factor grammar,’ in which only two syntactic classes were allowed, “topics, and authors” (Klatt, 1977, pp. 1345, 1353). It was designed around another model task, that of document storage and retrieval: in the rather topical examples given by a video presentation on the system, users could ask it to “give me the headlines,” “tell us about Nixon,” or “tell me all about China” (*Harpy*, 1971). The models which aimed for more general utility ended up by having lower accuracy.

Research on problem-solving also began to pursue greater practical usefulness, at the cost of generality and autonomy. This trend produced a strain of ‘expert systems’ research, covering fields from professional football, to cookery, to medical diagnostics. Most importantly for my purposes, these included several efforts which sought to effectively close the technoscientific loop, by partially automating the very process of scientific discovery. Like many others within the cognitivist movement, Joshua Lederberg was originally a researcher in another field (in his case medical genetics) who “quickly succumbed to the hacker syndrome” upon his early encounters with computers (Lederberg, 1987, p. 4). Though he arrived at Stanford intending to found a Department of Genetics within its medical school, he wound up working with Simon’s former student, Edward Feigenbaum, on a program called ‘Heuristic DENDRAL.’ This project unfolded over more than a decade, using the heuristic techniques pioneered by Newell and Simon along with the interactive computing and time-sharing promoted by McCarthy and Licklider, in an attempt to create software that could help to discriminate material compositions in mass spectrometry.

A mass spectrometer “is an instrument that converts molecules of a sample material into ions that are accelerated and measured one by one,” assigning numerical quantities to their masses. Determining the composition of the sample thus begins with a ‘change-making’ problem, determining—given known rules of chemistry—all the different mixtures of atoms which could

produce the masses observed (*ibid.*, p.6). This could be done by a ‘brute-force’ search through all the possibilities, but it was also a perfect candidate for heuristic programming: a variety of “contextual information” regarding valence and bonding could “be incorporated early into the combinatorics,” reducing a “blind generate-and-test search by several factors” (*ibid.*, p.7). In practice, the system would work by being supplied with a histogram from a mass spectrometer, and a coded standard description of a molecule. DENDRAL would generate graph-theoretic representations of the molecule, and in ‘dialogue’ with the machine, the chemist would use their knowledge of “likely places for breaks to occur” in the structure within a spectrometer; these would become ‘subgraphs’ which the chemist could alter and move, while the program would report back whether these hypothetical arrangements matched the masses recorded on the histogram (Feigenbaum in Lederberg, 1987, p. 22). Hence in cooperation with the machine, generating different likely candidate graphs and testing them against data from a mass spectrometer, the chemist would eventually find a structure consistent with the empirical data, more rapidly than by human reasoning alone.

This aspect of DENDRAL was quite successful in practical applications, and many types of mass-spectrometry software are now produced using these same principles, aiding chemists’ intuitions in laboratories worldwide. This reflects one of the persistent complaints of AI workers: now that such software is in wide use, and highly effective for its particular task, nobody really considers it ‘intelligent.’ The current systems, however, are more oriented toward brute-force and pattern-matching than heuristic programming, and Lederberg and Feigenbaum were already uncertain about the relevance of their project to the study of cognition. Was it just a useful tool, or a step toward understanding the nature of problem-solving and reason? In order to serve as the latter, they envisioned a different mode of operation for DENDRAL, in which it would be capable of

adaptively learning from the chemists which operated it, and eventually taking over much of the analytical process itself. The “ultimate fantasy” was a high-order DENDRAL system capable of parsing and learning from the published chemistry literature, attached directly to a spectrometer, and thus “learning directly from Nature” (Lederberg, 1987, p. 12). Perhaps as important as their actual output, the DENDRAL research team also epitomized the new techniques for computer-mediated scientific interaction and communication proposed by Licklider and others. Although there were occasional informal meetings, most of the group’s serious communication, whether from across the country or “a few yards down the hall,” was conducted by electronic mail. Thus, as Lederberg recalls, “innumerable proposals and drafts could be posted,” and run through scores of iterations, while “distance was no consideration, courtesy of the ARPANET” (*ibid.*, p.13). As Licklider had proposed, the facility for manipulating shared representations was a far more important feature of electronic communications media for scientific purposes than seeing the face of one’s collaborators.

Another ‘expert systems’ project along similar lines funded toward the end of ARPA’s institutional support for AI was ‘BACON,’ a system constructed at CMU that “discovers empirical laws” from sets of observational data using heuristics. From a table with measured temperatures, pressures, and volumes, for instance, it would be able to infer the ‘ideal gas law’ relating the quantities; it was also reportedly capable of rediscovering Ohm’s Law, Kepler’s Third Law, and myriad others (Langley, 1979). It applied simple heuristics, such as noting when two variables were trending in the same or opposite direction, or when a variable remained constant, then used that information to generate different abstracted, algebraic statements and test them against the data. BACON is most notable for giving rise, over a decade after its construction, to a direct collision between cognitivist AI and the sociology of scientific knowledge, with Peter Slezak’s argument that it is the “larger enterprise of cognitive science” which best brings into relief the “severe and

notorious inadequacies” of the Strong Programme first proposed by David Bloor (Slezak, 1991; cf. Bloor, 1976). As with GPS, the heuristics applied by BACON are quite general, abstract and indeed ‘weak’ as compared with our typical view of human insight. From this, Slezak concludes that it demonstrates a case of scientific discovery wholly divorced from ‘social factors,’ and thus a definitive falsification of externalism in the sociology of science.

If computers can re-discover the very same scientific laws as Kepler or Boyle, he argued, then their social circumstances cannot have played any causal role in the construction of said laws. While fascinating, I cannot explore the ensuing debate in detail (Fuller, 1989; Gorman, 1989; Slezak, 1992). Most significant for my purposes is the central premise in most responses to Slezak: that these programs are far from independent of human and social factors. Their ‘discoveries’ are “highly dependent on the way the data are structured by the programmer” (Gorman, 1989, p. 646), and actually show evidence for underdetermination. Given different presentations of the data, similar systems could in fact be induced to form ‘false,’ antiquated theories like the phlogiston model (Langley, 1987). The creation and programming of heuristics, along with the organization and input of data, are no less social processes for their apparently ‘common-sense’ character. Simon, unlike Slezak, recognized that imbuing a machine with this collectively-developed common sense was integral to the process of making it think, and making it a model for human thinking. Still more damningly for Slezak’s argument, by the time it was published, the weaknesses of the symbolic approach pioneered by Simon were growing increasingly evident, and institutional support had already substantially declined.

As argued by the later researcher and most famous current popularizer of AI, Ray Kurzweil, it turned out that the problems which we saw as most difficult and representative of the ‘higher’ faculties of human reason, like proving logical theorems or playing chess, were ‘easy,’ and “what

proved elusive were the skills that any five-year-old possesses” (Kurzweil, 1990, p. 70). Things like recognizing different objects in varying light, distinguishing spoken words and the breaks between them, and deploying common-sense knowledge about the world, proved to be far more challenging than the tasks adopted by early symbolic AI. Such setbacks gave rise to the well-known Heideggerian critiques of Hubert Dreyfus, first published starting in the mid-1960s, and culminating with his *What Computers Can't Do* (Dreyfus, 1979). Dreyfus' ideas served as a shibboleth for what would be christened in the 1980s as ‘Good Old-Fashioned AI’ (Haugeland, 1989). Those who held that the symbolic programme should be carried forward more or less as before, and that Dreyfus was utterly misguided, formed the camp of happily ‘Old-Fashioned’ cognitivists and AI researchers. Those who thought that perhaps he was on to something, at least occasionally, and that the techniques and systems employed for this research should change, forged an even more pluralistic bunch of new approaches.

More directly instrumental in the decline of symbolic AI was the UK government-commissioned Lighthill report. In his report, mathematician James Lighthill contended that

most workers in AI research and related fields confess to a pronounced feeling of disappointment in what has been achieved in the past twenty-five years. Workers entered the field around 1950, and even around 1960, with high hopes that are very far from having been realised in 1972. In no part of the field have the discoveries made so far produced the major impact that was then promised (Lighthill, 1973).

He likewise separated out multiple strands of the broader enterprise, using the rather convoluted ‘ABC’ mnemonic for ‘Advanced Automation,’ work in ‘Building Robots and Bridging’ the other two strands, and ‘Computer-based CNS research.’ Category A was where he was most skeptical, due to what he saw as the primary issue facing the entire enterprise, ‘combinatorial explosion.’ Whatever representations and heuristics we might devise that functioned well in restricted task domains, general intelligence required too large a database of knowledge, with potentially infinite possible

ways of grouping together and relating elements, all of which needed to be explicitly programmed in.⁴⁹ Where programs did achieve minor successes, programmers were often feeding information in very carefully-formatted ways that helped produce the desired answer (as with BACON). He singles out DENDRAL, by contrast, as the kind of ‘Category A’ research that did merit funding, where large volumes of knowledge were applied to specific task domains, supporting rather than replacing human reasoning. The other categories he saw as likewise involving a mixture of real achievements with excessive hype and ‘naïve optimism.’ Building computer models of whole nervous systems, or robots demonstrating human-level performance in simply picking up objects, for instance, was now increasingly recognized as a faraway dream.⁵⁰

Lighthill’s report has been described as calling “for a virtual halt to all AI research in Britain” (Crevier, 1993), but really he argued for more specialized, practical, and task-specific projects rather than ‘blue-sky’ research which aimed for general applicability or human-level performance. He closes the report by highlighting the ‘over-generalized euphoria’ which greeted Winograd’s SHRDLU, and characterizes it as a minor though notable success, yet emphasizes that we can only draw very limited conclusions from its successes:

In practice, a large computer together with very sophisticated programming using subtle new programming-language developments was found just sufficient to make slow conversation possible on the very limited material represented by the abstract table-top world; material restricted enough, for example, to allow resolution of ambiguities in natural-language sentences by classical theorem-proving techniques. Extension of the methods used to a much wider universe of discourse would be opposed violently by the combinatorial explosion (Lighthill, 1973).

⁴⁹ Optical character recognition gives a good example: as Lighthill noted at the time there had been little success in that regard, and though recognition of clear typewritten text is now widely available and reliable, it quickly fails if the base image is bad, and handwriting recognition is still quite unreliable. Humans, by contrast, are readily able to recognize letters across an enormous range of typefaces, handwriting styles, and sources of interference or distortion. We never need to explicitly be shown all of these variations, while symbolic AI programs apparently do—and the same goes for phonemes of spoken language.

⁵⁰ This observation, made with respect to the challenges inherent in a simple task like turning a knob situated amidst other items on a desk, opens Rumelhart and McClelland’s work on parallel processing (Rumelhart & McClelland, 1989, p. 5)

This combinatorial explosion was in one sense at the root of slow progress in AI. In another, though, it was simply an engineering problem. Contemporaneous with Lighthill's report, the US government was undertaking a shift in its funding priorities, and in 1972 ARPA became DARPA. After a decade of "significant funding" for AI and cognitive modeling under Licklider and his immediate successor, new management came in, "and people who weren't hostile to AI, but also weren't friendly, moved in." Funding wasn't immediately cut off, but gradually slowed, with the feeling that "AI had had more than its fair share and that other fields, especially the emerging area of supercomputing, deserved a chance" (Hendler, 2008, p. 2).⁵¹ Programs started in the early 1970s ran their course, and produced successes like SHRDLU and DENDRAL, still under ARPA funding, but in the "mid-to-late '80s...the lack of new research transitioning to industry made itself really felt—the seed corn had been consumed, and there was a dearth of new crop" (Hendler, 2008, p. 3). Because of Lighthill's report, I highlight 1974 as the year when the heyday of cognitivism came to a close, and the ensuing 'AI winter' marked an interim period of reduced funding for cognitive science as well. Equally significant though in precipitating the decline was the 1969 Mansfield Amendment, which limited US Department of Defense spending to projects with direct military applicability. AI had never really been 'big science,' but it was well on its way then to becoming considerably smaller. The coining and professionalization of 'cognitive science' in the late 1970s was thus in part a pivot away from the dream of autonomous thinking machines, back toward greater emphasis on the use of computing as a tool for understanding the mind and brain.

⁵¹ This turn to supercomputing meant that conventional heuristic, symbolic AI programming was facing competition on two fronts, which I discuss in more detail below: from the alternative, neoconnectionist and embodied cognition models, but also from the school of programming which sought to address combinatorial explosion by throwing exponentially more computing power at the problem, and skipping the heuristics altogether.

It was cognitivism which sought most fervently to make the root metaphor of cognitive science into a literal expression: defining cognition not simply as a physical phenomenon which could be modelled with machines, as Turing and the cyberneticians had, but as identical with computation itself. In the process, it learned a good deal about processes which may underpin higher-level symbolic reasoning in humans, reconceptualizing cognition in general according to an understanding of how the cognition of scientists and mathematicians operated. Programs like DENDRAL and Logic Theorist could far outstrip the capacities of many humans in their respective task domains. Yet this only highlighted the twin deficiencies of this strategy for modelling cognition: it was less well-suited to real-world perceptual tasks, and despite having enormous stores of information, these programs seemed to lack anything resembling tacit knowledge or common sense. They could not seem to form concepts based on implicit cues, as we do so readily—for instance simply keeping track of pronoun references in a free-flowing conversation—and in what amounts to the same problem of tacit knowledge in a more ‘embodied’ form, they could not make inductions about how things would happen in the real world, from a lifetime of experience with its physics.

While rarely rejecting outright the notion that symbolic computation plays some role in human cognition, researchers after the 1980s began recognizing more widely that some of the analog approaches of the cybernetic era might be more fruitful in addressing these problems. At the same time, partisans of symbolic AI still contend that the early predictions were valid, despite their over-optimistic timelines, and have continued working to expand their databases and heuristics to sufficient size for encompassing common-sense understanding of the world.⁵² The achievements of these projects, while impressive in some respects, ultimately pale in comparison to the achievements

⁵² Douglas Lenat’s Cyc project is still active as both an open-source project and commercial venture, receiving periodic media coverage in which its creator’s optimism for symbols and search is unflagging though the system still fails to demonstrate ‘real’ intelligence and common sense (“The word: Common sense,” 2006).

of another now-ubiquitous form of hybridized human-computer intelligence, as I explore in the next section. Natural-language conversation with computers, another longstanding goal of this approach, has also proven elusive, with entrants in the Loebner prize competition—an annual contest modelled on Turing’s imitation game—periodically fooling a few judges, but never coming close to systematically indistinguishable from human interlocutors. Some efforts in this contest have placed greater emphasis on actually storing some tacit knowledge about the world, whereas others cultivate a deliberately unusual verbal ‘personality’ to make any seeming variations from conversational norms seem more natural.⁵³ Yet all seem susceptible to the critique that they don’t demonstrate *real* intelligence, simply a complicated set of rules for transforming symbolic expressions (in its broadest form, this is Searle’s ‘Chinese Room’ argument).

Cognitivist researchers rail against this tendency to think that ‘if it works it isn’t AI’ (McCorduck, 2004, p. 204): we imagine AI as something faraway and mysterious, but when something once thought to be solely achievable by human intelligence is turned into a computer program, for some reason the mystery disappears and we no longer see it as intelligent. There is some merit to this critique, particularly when it comes to sweeping arguments against the very idea of AI. Yet there is also an element of duplicity, inasmuch as what *has* worked has been very different from what the early practitioners proposed. The first great success story of IBM’s return to the field, its Deep Blue chess-playing computer, relied on brute-force search rather than advanced heuristics, evaluating 200 million potential board positions each second (Campbell, Hoane Jr., & Hsu, 2002)

⁵³ The most famous example of a program which used this ‘loophole’ to successfully convince a number of human judges of its humanity was ‘PARRY,’ a program which long predated the Loebner competition. Written in LISP, it was designed to simulate the responses of a paranoid schizophrenic: lashing out in anger, turning the conversation toward his delusions about the Mafia, and so on (Colby, Weber, & Hilf, 1971). This is a good case in point for what Harry Collins calls ‘repair, attribution, and all that,’ and how humans can be primed to sometimes ‘repair’ another’s conversational contributions more heavily, whereas if a program designed to simulate a more ‘normal’ conversation slips up and delivers a non sequitur response, we’re more inclined to read it as evidence for its inhumanity (Collins, 1990).

instead of incorporating any profound insights into human expertise. Newell and Simon, by contrast, equated the heuristic approach with AI itself, seeing brute-force techniques instead as simple computing. Indeed, if we accept Deep Blue as a form of AI, it is of a type which severs the longstanding ties between AI and cognitive modelling: it simulates some specific capacities associated with human intelligence, but in a way that could never be plausibly ‘implemented’ in our brains.⁵⁴ Other strands which I explore below sought instead to reaffirm and strengthen the ties between AI, the modeling of minds, and the study of brains.

What has been christened now as ‘GOFAI’ retains an intrinsic appeal. At its core, it seeks to take what we know about the mind from psychological experiment and disciplined introspection, translating it into programs that no longer exactly model ‘natural’ human cognition, but rather implement cognition as such and in general. Cognition is thus constituted through symbolic representations and laws for manipulating them—a ‘language of thought’ (Fodor, 2008)—which may have some particular quirks whether implemented in a digital computer or a brain, but it is nevertheless one unified phenomenon. As the current label indicates, in its pure form this now has a whiff of the antiquated, but it remains a potent thesis. As many of its actors attested, it was bound together less by any theoretical unity than by a shared set of tools: programs, programming languages, and digital computers. In the next section, I explore some of the ways in which the field has continued to rely on these tools, the former becoming still more pluralistic as the latter have continued to transform. Cognitivism became ‘old-fashioned’ as the modeling of mind returned in various ways to neurobiology, embodiment, and analog computing.

⁵⁴ IBM’s later Watson project presents a more complicated case of hybrid human-computer intelligence leveraging the Internet and Wikipedia, which I discuss in the next section.

3.3: Post-Cognitivism: PDP, Connectionism, and a return to neurobiology

The theory of cognitivism, in its literal equation of mind and computation through multiple realizability, was profoundly influenced by the adoption of a particular class of computing systems (as Newell himself contends: (1982, p. 13). The implications of this engagement with digital computing were felt across the disciplines which constituted cognitive science. Many of the foundational techniques for the wider field of computer science were first forged within the same circles, including those for time-sharing, higher-level programming languages, and database manipulation. Figures like McCarthy, Newell, Simon and Minsky were the most closely engaged with the modeling tools and practical goals of AI, while others like Miller, Bruner, and Ulric Neisser focused more on reframing psychology in light of cognitivist ideas about information-processing. In the 1980s, however, interest started to turn back toward the same classes of distributed, parallel, and ‘continuous’ systems previously pursued by cybernetics. The work of the Parallel Distributed Processing (PDP)⁵⁵ group reinvigorated the neural-network approach from the Perceptron era as a method of cognitive modeling, now to be understood not as a psychological offshoot of cybernetics but an alternative toolkit for cognitive science. At the same time, the development of new tools for computer-aided neuroimaging drove renewed interest in and institutional prioritization of the brain as an obligatory passage point for understanding human psychology. This final section briefly chronicles these developments, and sums up my argument on mediation in cognitive science, past and present.

As proposed by a core of researchers centred on the Institute for Cognitive Science at the University of California, San Diego, led by David Rumelhart and James McClelland, PDP marked a

⁵⁵ Not to be confused with the DEC hardware having the same acronym.

return to ‘physiologically-flavoured’ modelling (Rumelhart & McClelland, 1989, p. 21). It also marked a new geographic centre for cognitive science in California: first at UCSD, then at Stanford, though McClelland also spent much of his career at CMU. The PDP group argued that humans are ‘smarter’ than computers not in speed or precision, but in the facility with which we handle perceiving objects in natural scenes, understanding language and retrieving contextually appropriate information, and “a wide range of other natural cognitive tasks.” We have this advantage because “the brain employs a basic computational architecture” (Rumelhart & McClelland, 1989, p. 14) that is better suited to such tasks than the programs of the cognitivists. PDP models shared a structure “inspired by basic properties of the neural hardware,” consisting of “a large number of highly interconnected elements... which apparently send very simple excitatory and inhibitory messages to each other and update their excitations on the basis of these simple messages” (*ibid.*, p.21). This was not a matter of wholly rejecting more psychologically-flavoured and sequential, symbolic models, but of shifting focus and exploring a category of systems which seemed more plausible to the group as accounts of the ‘microstructure’ of human cognition. Carrying on and reinterpreting the hardware/software metaphor of multiple realizability, they argued our biological hardware was simply ‘too sluggish’ for serial processing models to serve as accounts of this microstructure (*ibid.*, p.23). Instead their models were composed of many parallel units processing simultaneously, with the computational work being carried out in the ‘spreading activation’ across their interconnections: hence the term ‘connectionism,’ coined by Rosenblatt, was reapplied to this new approach. It was explicitly framed as a collection of “alternatives to the models that have dominated cognitive psychology for the past decade or so” (*ibid.*, p.xi).

The PDP group was influenced by the work of David Marr on vision and differing cognitive ‘levels of description’ (Marr & Poggio, 1976), as well as experience with the DARPA-funded

HEARSAY speech understanding project (Erman, Hayes-Roth, Lesser, & Reddy, 1980; Rumelhart & McClelland, 1989, p. 43). Alongside the Perceptron, the other most significant precursor from the cybernetic era was Selfridge's Pandemonium model (Rumelhart & McClelland, 1989, p. 42), which used layers of 'yelling demons' to simulate the process of visual perception.⁵⁶ The models of PDP were quite similar to the Perceptron in their basic functioning, involving many interconnected nodes with dynamic weights determining whether or not signals are passed on. The revitalization of connectionism is often described as owing to the construction of networks with 'hidden' layers of nodes, as opposed to the single layer Perceptrons; while these hidden layers did allow the PDP models to accomplish more, including the 'XOR' function and determinations of connectedness that Minsky and Papert had shown to be beyond the capacities of Perceptrons, this was already well understood. What had been absent hitherto were good techniques for adjusting their connection weights and thus 'training' networks with these hidden layers. Though some tasks were amenable to 'unsupervised' adaptation with simple rules, derived for instance from Donald Hebb's research—strengthening connections between nodes that tended to activate together—making a neural network capable of character recognition and other complex tasks required active training. A considerable portion of the two volumes of *Parallel Distributed Processing* was occupied with mathematical functions for training networks and adjusting connection weights, propagating error signals back through the network when outputs did not match the desired results (Rumelhart & McClelland, 1989, p. 133).⁵⁷

⁵⁶ Selfridge, recall, was Wiener's research assistant, and it was a presentation he gave at RAND on his chess-playing program which first inspired Allen Newell's interest in AI (Crevier, 1993): hence he was an important figure in the development of both prior eras of cognitive science, but cognitivism was less interested in pursuing his specific parallel-processing approach.

⁵⁷ Many of these and later innovations in training techniques and 'learning algorithms' for neural networks are owing to the work of their collaborators Geoffrey Hinton and Terence Sejnowski (Hinton, 2007).

The PDP approach was envisioned as an ‘exploration,’ a plurality of models and a collection of tools rather than a systematic theory. It embodied certain general principles for understanding cognition, with varying details according to particular intended applications—be they practical tasks or models of mind (Rumelhart & McClelland, 1989, p. 145). Their most important shared feature, the correlate of their parallel architecture, was their distributed memory: for symbolic models, “knowledge is stored as a static copy of a pattern,” and retrieving it amounts to copying it from long-term storage into “a buffer or working memory;” in connectionist models, by contrast, “the patterns themselves are not stored,” but rather the connection strengths between units allow them to be re-created (*ibid.*, p.42). Along with serving more readily as models for adaptation and learning, they were likewise more physiologically plausible in their capacity for ‘graceful degradation’ following so-called ‘lesioning’ (*ibid.*, p.134). While a conventional computer abruptly fails to operate if one of its key hardware or software components goes missing, nodes can be removed from a neural network with only minor decreases in task reliability. These gradually worsen as more nodes are removed, similar to the effects of neurodegenerative syndromes and surgical brain lesioning in biological organisms.⁵⁸ And unlike the discrete symbolic representations of digital computers, computation in these networks is carried out by the continuously variable weightings of connections between nodes. Hence these are once again analog machines in an important sense, just as the machines of cybernetics were. Yet the story of post-cognitivist efforts to model the mind is equally one in which these long-standing divisions are growing increasingly blurry.

PDP models, to begin with, are not machines themselves but mathematical abstractions—directed graphs of interconnected functions—typically implemented as simulations on increasingly

⁵⁸ This phenomenon of graceful degradation is distinct from the problem of engineering reliable digital machines from unreliable components, as addressed by von Neumann (Von Neumann, 1956), though as in his other writings and those of the broader cybernetics group, that paper is also concerned with the commonalities between machines and organisms.

commonplace, low-cost digital computers. The rise of this new connectionism was contemporary with that of personal computing, and so there was little incentive to build specialized devices as Ashby and Rosenblatt had. Hence even if the functioning and architecture of the models was continuous and parallel, critics could contest that they were at base still carried out on digital hardware; conversely, advocates of PDP argued that an 'arbitrary computational machine' could be put together out of parallel-processing units, including a Turing-equivalent machine, with the true limitation being that "real biological systems cannot be Turing machines because they have finite hardware." But as they noted, "with external memory aids (such as paper and pencil and a notational system) such limitations can be overcome as well" (Rumelhart & McClelland, 1989, p. 119).

Following their 'microstructural' turn, connectionists could contend that it was most plausible to understand human cognition as akin to connectionist models on a subsymbolic, neurophysiological level, and the challenge was then to understand how these served to implement higher-level symbolic processing of the type emphasized by earlier AI research (Smolensky, 1988; cf. Fodor & Pylyshyn, 1988). Some cognitive scientists and philosophers of mind would instead offer stronger 'eliminativist' arguments with respect to our "common-sense conception of psychological phenomena," including the very notion of 'mind,' and suggest that connectionist modeling and brain imaging would replace these 'radically false' theories with a completely new account (Churchland, 1981). In this sense cognitive science has come to once again incorporate the view that originally it attacked, that mind is a non-scientific fiction. Such eliminativism is only one minority viewpoint within the big tent of cognitive science, however, and was long ruled out by the cognitivist vision of mind as a distinct 'software' level of description. It regains appeal only with the return of connectionist models, and new techniques for rendering visible the biological brain. Yet it is more commonly recognized that "successful lower-level theories" like connectionism "generally serve not

to replace higher-level ones, but to enrich them,” explaining their successes and failures, filling in gaps and unifying “disparate higher-level accounts” (Smolensky, 1988, p. 22).⁵⁹

Critics of reductionism, moreover, often take issue with an assumed bio-somatic ‘fatalism’ implicit in the idea, be it a reduction to neuroscience, to computation, or to both. To suggest that the human mind is nothing more than patterns of spreading activation in nerves, or a thing computable by a Turing machine, or both, is often seen as an attempt to impugn the creativity, free will, even the reality of the human mind as it has long been understood. At times it is, either openly or implicitly.⁶⁰ More often, however, the valence of the analogy seems to run the other way, an elevation of the capacities of machines to those of human minds: the intuition that electronic switching systems, be they biological or artificial, have sufficiently universal capacities to account for the full spectrum of adaptive behaviour, intelligent thought, and conscious experience. Though not universally held or even consistent in its manifestations, this intuition has long been closely tied with another: that logic and mathematics represent the ideal mode of description for both sets of phenomena, electronic computation and biological cognition. This intimate entanglement is embodied in Douglas Hofstadter, another important figure in the post-cognitivist, ‘AI winter’ era. His reflections on the techniques of symbolic ‘recoding’ and recursive level-crossing involved in

⁵⁹ Bruno Latour draws a similar conclusion of more general import, remarking on how Paul Churchland, his “former colleague in San Diego,” carries in his wallet a printed brain scan of his wife (and frequent collaborator) Patricia, insisting adamantly that before long we will all think this way, “recognizing the inner shapes of the brain structure with a more loving gaze than noses, skins and eyes!” Paul’s eliminativism holds that “once we have a way of grasping the primary qualities”—on his view, the structure of neural connections, but perhaps more fundamentally the microstructure of neurons themselves or the biophysics of DNA or the “information content of the whole body measured in gigabytes”—then “we can eliminate as irrelevant all other versions of what it is to be a body, that is, to be somebody.” Latour proposes instead that instead of accepting Churchland at his word “as the reductionist and eliminativist he claims to be,” we recognize his attempted reductionism as “adding one more contrast, one more articulation to what it is to have a body” (Latour, 2004a, pp. 224–225).

⁶⁰ The Churchlands’ writings offer ample evidence for this reading of connectionism and neuroscience. Though I cannot explore this contrast in any detail here, wholesale reductionism seems to be more commonly espoused in popular-neuroscience literature than it is amongst the broader community of cognitive scientists (Eagleman, 2011; S. Harris, 2012).

Gödel's incompleteness proofs led to a massive, influential tome (Hofstadter, 1979), expounding the view that our capacities for recursion and analogy-making constitute the core of cognition, and that the best way to understand the mind is to construct modest mathematical and computer models of those specific processes. As a recent journalistic account observes, Hofstadter's book "arrived on the scene at an inflection point in AI's history," when "ambitious AI research had acquired a bad reputation" (Somers, 2013). The 'winter' had started to set in, and so Hofstadter's proposals were influential in some quarters, singled out for instance by Rumelhart & McClelland, but in other respects he found himself diverging from the mainstream. His 'romantic dream of AI' positioned him against the "intolerable amount of hype... self-promotion, grandstanding and implausibly grandiose claims" of an earlier generation, which likewise sought to understand human intelligence, but also from the development of a new results-oriented approach to AI which had largely set aside that goal (Ekbia, 2008, p. 12).

Despite mutual sympathies, the approach he pursued with his Fluid Analogies Research Group also diverged




Figure 20 - Copycat table of results

from the PDP approach in not aiming for a high degree of physiological plausibility. Instead, their various models sought to model analogy-making and concept-formation—which Hofstadter saw as one and the same process—in 'austere micro-worlds' that nevertheless remained intuitively comprehensible as analogous to human cognition. The best example of these is Copycat, created by Melanie Mitchell and Hofstadter, a program which was neither "symbolic nor connectionist, nor was it intended to be a hybrid of the two (although some might see it that way); rather the program has a novel type of architecture situated somewhere in between these extremes" (Hofstadter and Mitchell

in Hofstadter, 1996, p. 205). It had what they called an ‘emergent architecture,’ in which behaviour was a “statistical consequence of myriad small computational actions,” using “statistically emergent active symbols” (*ibid.*). The task domain for Copycat was to discover analogies between transformations of short letter strings in a psychologically plausible way. So it would be shown the transformation “AABC -> AABD” and asked to copy the pattern using the input string ‘IJKK;’ its output came from the aggregation of many ‘runs’ as depicted in Figure 20 (*ibid.*, p. 238). Most people considering this problem, they say, would quickly see that the ‘rule’ embodied in the first change was to replace the rightmost letter by its successor in the alphabet; the question is then whether the rule should be “transported rigidly” to the other string, or whether the two ‘k’s should be treated as a unit, and both changed to L. Just as most people find ‘IJLL’ the most intuitive analogy, but some would argue for ‘IJKL’ or ‘JJKK,’ Copycat arrives at the same possible answers using a version of ‘mental fluidity’ which its creators contend is faithful to our own.

The program uses an architecture based to some degree on the earlier HEARSAY speech-understanding project, involving levels described as a ‘Slipnet,’ the long-term memory and “site of all *permanent Platonic concepts;*” a ‘Workspace,’ or “locus of *perceptual activity;*” and a ‘Coderack,’ or “stochastic waiting room” in which “small agents that wish to carry out tasks in the Workspace wait to be called” (*ibid.*, p. 211, emphasis original). These ‘small agents’ are themselves composed of symbolic computations, as proposed by the cognitivists, but in other crucial ways Copycat more closely resembles cybernetic and connectionist models. The coordination of different ‘scout’ and ‘effector’ codelets, which respectively look at possible transformations to “evaluate their promise,” and create or destroy candidate structures in the Workspace, involves a considerable degree of randomness, which FARG sees as the key to ‘fluidity.’ Like several of their other models, it combines ‘bottom-up’ or unfocused searching agents with ‘top-down,’ more strongly goal-directed

ones, in a variable noted on the graph as ‘temperature,’ a ‘regulator of open-mindedness’ (*ibid.*, p.228) which changes over successive runs. The system starts by exploring possibilities in a relatively random fashion, but gradually decreases in temperature as more parallel agents start to focus around the same answer, or occasionally increases if a new candidate recurs. This process of dynamic competition and random, generative ‘perceptual objects’ was presented as essential to both human cognition and AI, and despite its simplicity, as likely more fruitful in both domains than the purely symbolic, sequential architectures which produced programs like Logic Theorist. This view, that human cognition was constituted and best modeled by hybrid systems—with both analog and digital, serial and parallel elements—had been expressed long before by von Neumann, but received renewed attention following the 1980s (Von Neumann, 2012; Minsky, 1986). Though small in scale compared to some of the other projects discussed in this section, Hofstadter and FARG produced some of the most interesting explorations along these lines. And like earlier approaches to cognitive science, they were in some respects using the “scientist as a model of human nature,” and scientific thinking as their image of thought in general (Cohen-Cole, 2005). They retained a focus on the formation of inferences, hypothesis, and categories, but added the intuition that at the root of all these phenomena lies analogy and metaphor. Whether this is sufficient basis for a scientific account of cognition as such remains an open question, but as I have argued, analogy has undoubtedly played a central role in the history of cognitive science.

Alongside this new conception of how psychologically plausible systems might be constructed within micro-worlds, a different side of cybernetics received renewed attention: that of ‘lifelike’ robots operating in the real world, directly inspired by the work of Grey Walter. This movement also emerged from the MIT AI laboratory, but rejected its earlier focus on symbolic representations and specialized subproblems, arguing instead that for the ‘harder’ problems of basic

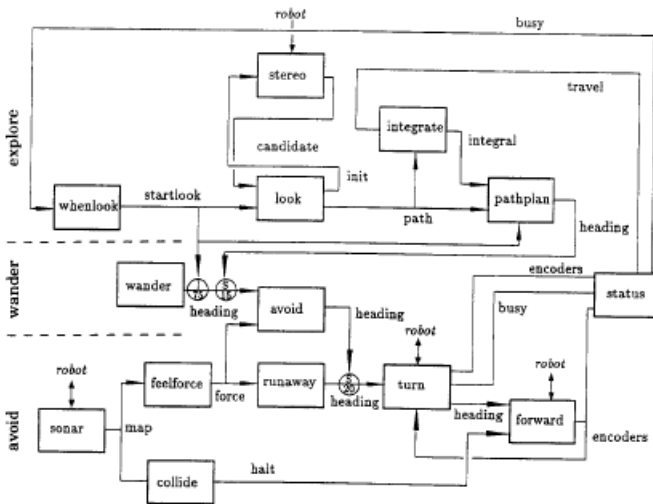


Figure 21 – Rodney Brooks' subsumption architecture

intelligence and real-world adaptive behaviour, “explicit representations and models of the world simply get in the way”. It might be better to simply “use the world as its own model” (Brooks, 1991, p. 140), dynamically gathering information as required. Rodney Brooks' ‘subsumption architecture’ (Figure 21) involves feedback

connections akin to those of cybernetic robots, where each layer and module is semi-independent, sending and receiving messages to one another without a central control, but now implemented in transistors and integrated circuits. Two of the AI Lab robots built according to this architecture were named Allen and Herbert (after Newell and Simon, of course). Each had distinct layers allowing first for basic obstacle avoidance, another leading it to wander randomly, and a top layer directing it to explore new areas and determine paths (Brooks, 1991, p. 152). These higher-level goals were not central governing representations, however, and the lower-level object avoidance function in no way depended upon them. Like Hofstadter's approach this was parallel and synchronous, involving a significant element of randomness, and at the same time distinguished itself from connectionism; but quite unlike Hofstadter, Brooks argued that working directly in the real world, with input systems ‘perceiving’ and output ‘effectors’ acting upon it, was the only way to proceed for AI and cognitive simulation, and “anything less provides a candidate with which we can delude ourselves” (*ibid.*, p.140). Brooks went on to become one of the founders of iRobot, a corporation responsible for the highly successful Roomba vacuum cleaner, and a beneficiary of DARPA's recent return to funding AI research. The most visible face of military robotics to date has been the rather distinct

field of remotely-operated ‘dumb’ robots like aerial drones and bomb-disposal units, but iRobot is one of many firms pursuing more autonomous robotic systems, capable of manipulating objects, driving cars, or quadrupedal locomotion with only ‘high-level direction’ by humans.

All of this raises evident ethical issues, particularly as we consider the possibility of autonomous robots fitted with weaponry. It marks not only a return to the ideas of cybernetics, but their return to the fold of military funding and military interests. This aspect of the contemporary cognitive sciences should in no way be downplayed, and as I discuss in the remainder of this section, DARPA is once again a primary institutional supporter of research in the field, from robotics to ‘neuromorphic computing,’ and ‘closed-loop’ technologies for simultaneously representing brains and intervening in them therapeutically (Defense Sciences Office, 2013). Yet I have equally tried to emphasize the heterogeneity across this field’s history, and the ways in which cybernetic and cognitivist discourses have diverged from their military origins and applications. The closed-world discourse of Cold War science gave rise to a conception of the human mind as an abstract computational structure, composed of symbols, heuristics, and rules – as Edwards has contended and others have supported. These could be realized either as lines of computer code, or within organisms as plans: “*any hierarchical process in the organism that can control the order in which a sequence of operations is to be performed... essentially the same as a program for a computer*” (Miller et al., 1960, p. 17). Yet no less a part of cognitive science, and no less entangled with military and industrial research programs, we find the competing hypothesis that such plans were merely a retrospective way of communicating about goal-directed behaviour, and that a better way to model cognition was as a continuous process of situated, embodied interaction with the environment (Agre & Chapman, 1990; Suchman, 1988). This alternative account had ties with non-representational robotics, parallel-processing, and theories of cognition as a distributed and extended phenomenon, drawing these

together to reimagine the mind: from a closed system operating within a brain or computer according to heuristic rules, to an open system cutting across boundaries of brain, body, and world. Work within this stream remained computationally-oriented to varying degrees, bringing its own distinctive models and programs.⁶¹ These employed parallel processing and embodied action to simplify and speed up the computations needed to achieve tasks, addressing the combinatorial explosion confronting the symbolic approach.

This new mode of non-cognitivist cognitive science was deeply influenced by anthropological research, and some of its core works have become important touchstones for science and technology studies as well (Suchman, 1988; Hutchins, 1996). A related movement grew out of what has been called ‘third-wave’ cybernetics, the autopoiesis theory of Francisco Varela and Humberto Maturana, which became an ‘enactive’ account of cognition (Maturana & Varela, 1980; Varela et al., 1991). “One of the first statements” of this viewpoint (Vera & Simon, 1993, p. 12) gives a good indication of the subtle shift in its stance toward computer models: “The computer, like any other medium, must be understood in the context of communication and the larger network of equipment and practices in which it is situated” (Winograd & Flores, 1986, p. 5).⁶² Here we return to

⁶¹ Agre and Chapman produced Pengi, a program designed to operate without central plans or representations, instead through “a continual, participatory interaction with its environment,” which was a game they called ‘Pengo’ (Agre & Chapman, 1990, p. 24). Edwin Hutchins’ *Cognition in the Wild*, as already mentioned in previous chapters, is another important work here, characterizing cognition as a distributed and mediated computational process, while Lucy Suchman’s models in the formative period for her work were ‘conversational’ interfaces developed at Xerox PARC (Hutchins, 1996; Suchman, 1988).

⁶² Winograd, the creator of the symbolic AI program SHRDLU and its ‘block world,’ had by this point abandoned the ‘GOFAI’ approach and begun working at Xerox PARC, then Stanford; influenced by the phenomenological arguments of Hubert Dreyfus and his coauthor Fernando Flores, as well as his likeminded colleagues at PARC, he had become a proponent of connectionism, situated action, and other alternative approaches. His work, like Suchman’s, continued to focus on human-computer interaction, and as he notes there was a continuing influence of Licklider and his IPTO in the framing of this research: Licklider “was into human-computer interaction before he got into computers; that was his background... IPTO hype, of course, is always a question of what people think will get funded or how they should paint it so it will get funded. Certainly, within the AI work, and computer science as a whole, of course, is a much bigger question. There was a taken for granted assumption, which is if you can get computers to be more like people they’ll interact better” (Winograd, 1991, pp. 45–46)

the sense in which STS came to recognize that ‘cognitive explanations’ need not stand in opposition to alternative sociological and otherwise ‘externalist’ accounts (Keller et al., 1996). As the Slezak controversy suggested, and as Harry Collins has explored at length, there is indeed a tension between ‘GOFAI’-style symbolic accounts of cognition—often specifically addressed to the process of scientific discovery—and those favoured within STS. The extent of the opposition even in this case remains debatable, particularly given the organizational orientation of Newell and Simon’s research program. But many later developments in cognitive science sought to further emphasize the situatedness of cognition within a network of bodies, materials, media, and social practices. Far from representing a new consensus, these new accounts are heterogeneous, and debates persist over not only which view is preferable, but to what extent they are even opposed. Herbert Simon and his allies argued that situated action was an insufficient replacement for plans and programs, susceptible to better understanding in symbolic terms, while those advocating for situated action contended that they were neither one monolithic group, nor did their pluralistic approach aspire to wholly replace cognitivist ideas (Vera & Simon, 1993; Agre, 1993; Suchman, 1993).

The same work which inspired Latour to recant on ‘cognitive explanations’ seems to furnish the best way to understand this dichotomy. As Edwin Hutchins argued, following a lengthy participant observation aboard the *Palau*, a U.S. Navy vessel, the problem of navigating on the ocean can be solved in two ways, both best understood as distributed cognitive systems. The approach taken by the *Palau*’s crew relied on many human actors coordinating the use of sighting instruments, charts, shipboard telephones, and other technologies in an intricate computational process, the “*propagation of representational state across a series of representational media*” to determine the vessel’s

position and steer it safely (Hutchins, 1996, p. 117).⁶³ This relies on explicit symbolic representations and a great deal of carefully engineered instrumentation, culminating in official matters of fact: knowledge of the vessel's position on a map and a course plan for reaching the intended destination. Hutchins cites Simon's theories as applying "very nicely" to this approach (*ibid.*), a form of sequential processing broadly compatible with the cognitivist image of thought, and similar to the kinds of processes Simon studied with Kennedy, Chapman, and Newell at RAND.

Yet he argues at the same time that this mode of thought is "historically contingent," and that the same problem of long-distance navigation on the open ocean was solved by the navigators of the Pacific Basin using an entirely different approach, leveraging situated, embodied action. No less effective or in some sense 'computational,' it relied not on technological instrumentation but the mnemonic technique of representing the night sky in terms of 'linear constellations,' steering courses between widely-separated islands by continuously tracking their position and heading with reference to those constellations and a third 'reference island,' sometimes imaginary (Hutchins, 1996, pp. 65–93). This more closely resembles the practice of 'pre-modern' European navigation as well, before the wide availability of key instruments like accurate compasses, charts, and the tables of logarithms which so interested Babbage. Though evidently also involving representations of some kind, and orally-transmitted media like the system of constellations, this approach only ever culminates in situated knowledge, not a 'god's eye' conception of the vessel's position on a map, nor an explicit course plan.⁶⁴ This perspective on situated action, as explored by Hutchins and Suchman among

⁶³ This process, incidentally, has been greatly transformed since Hutchins' research by the implementation of the Global Positioning System of satellites, as he fully anticipated at the time: *ibid.*, p.33.

⁶⁴ Ethnographic evidence shows that Micronesian navigators have great difficulty indicating their position on a Western-style map, but nevertheless throughout a voyage can at any time "accurately indicate the bearings of the port of departure, the destination, and other islands off to the side of the course being steered, even though all of these may be over the horizon and out of sight" (Hutchins, 1996, p. 66). Few, if any participants in the distributed cognitive system on a modern naval vessel could pull this off.

others, has proven influential within both contemporary cognitive science (Clark, 2003; Noë, 2004) and STS (Keller et al., 1996; Alač & Hutchins, 2004).

Close attention to the original debates shows that we should not view it as mutually exclusive with cognitivism, however, or as a crusade against the vision of humans as isolated information-processors. Rather, each constitutes a distinct and historically significant way of organizing and mediating cognition, each integrating social and computational dimensions in its own fashion.⁶⁵ Above all, what I wish to emphasize is the thread uniting even seemingly disparate positions like situated action and cognitivism with the enterprise of cognitive science: the analogy of mind as computation, however differently it may be interpreted and implemented. This implementation is labour-intensive and imperfect, involving a great deal of extension of our own cognition through interaction with external media. This is a major focus in Hutchins' cognitive ethnography, and it is given perhaps its most systematic treatment by Andy Clark, who contends that through its native plasticity, the human brain is "poised for profound mergers with the surrounding web of symbols, culture, and technology" (Clark, 2003, p. 208). This view that we are 'natural born cyborgs,' always-already positioned by evolution to extend our capacities for action and thought to form extended cognitive systems, informs this analysis throughout.⁶⁶ Whatever effects media may have within

⁶⁵ The site of Hutchins' research also indicates that the situated action perspective is no less affiliated with the apparatus of military operational research than was cybernetics or cognitivism. Another useful definition of the related concept of embodied cognition is as follows: "Embodiment is the surprisingly radical hypothesis that the brain is not the sole cognitive resource we have available to us to solve problems. Our bodies and their perceptually guided motions through the world do much of the work required to achieve our goals, replacing the need for complex internal mental representations. This simple fact utterly changes our idea of what 'cognition' involves, and thus embodiment is not simply another factor acting on an otherwise disembodied cognitive processes" (Wilson & Golonka, 2013, p. 1). Embodied cognition is not an alternative to the idea of mental computation, but a way of rethinking how computation in the brain works, and how taking into account the resources of our active, embodied situation may greatly simplify our cognitive models (Clark, 1998; Noë, 2004).

⁶⁶ Clark's work in turn is inspired by writings of Haraway and other 'postmodernist' thinkers as well as the original formulation of the cyborg-concept by Clynes and Kline, as discussed above in this text. While he is not the only one to express this view by any means, and does not bring any especially novel evidence to bear on the question, I find his contention that "human thought and reason is born out of looping interactions between material brains, material bodies, and complex cultural and technological environments" (Clark, 2003, p.22) well-framed and persuasive.

science and for the masses, this is not a pernicious outside influence befalling the ‘natural’ mind, nor a novel phenomenon of the twentieth century. Rather, that which is most distinctive about human nature is our capacity for mediating our cognition, through shared patterns and artificial constructions. In this sense, the development of computing machines and the analogy of mind itself with computation culminates a journey of millennia which began with the diffusion of language and writing.

From the analog machines of cybernetics, to the index cards Simon used to have his family simulate Logic Theorist, on through the languages, programs and digital machines of cognitivism, and the pluralistic turn back to biology, embodiment, and the ‘cyborg’ in the 21st century, exploration of this analogy has been the hallmark of cognitive science. Other fields such as psychology, like Miller recounted at the outset of this section, may find it useful in their own ways to make common cause within this interdisciplinary collective, often for funding purposes and particularly by incorporating the metaphors and tools of information-processing (Miller, 2003). What distinguishes the cognitive subfields of various sciences – cognitive psychology, cognitive neuroscience, and as I discuss below now even ‘cognitive computing’ – from the core figures of the cognitive science tradition is that the former employ the computer as a mediating device to access some other domain: be it the laboratory subject, the brain and CNS, or simply more effective applications for computers. Those most properly seen as cognitive scientists, beyond mere self-identification, are those most fascinated by the process of modeling and mediation in itself. All of the examples I have selected lie somewhere near this hypothetical core of cognitive science, but perhaps the most emblematic figures for each era are Warren McCulloch, Herbert Simon, and Douglas Hofstadter. Their writings are shaped by the technological and cultural contexts of each period, but they share the fundamental drive to understand the essence of our selves through

mathematics and computation. The traditional priorities of whatever disciplinary circles they travelled in were of secondary importance to this goal. For this very reason, however, they proved instrumental in the processes of cross-disciplinary exchange by which the various subdisciplines underwent their own forms of ‘cognitive revolution.’

Before concluding this section, I wish to examine some even more recent developments, considering the interplay between the trends I have outlined and the current interest in ‘brainhood’ as a mode of understanding subjectivity (Vidal, 2009). Neurobiology has come full circle, from a core interest of the cybernetics group, to its near-elision by cognitivism, to the present day, when more than any of the other subfields making up the cognitive sciences it stands as the obligatory passage point to understanding the self. The ‘neuro’ turn is one which has received more sustained critical attention than the ‘cognitive.’ Many have argued that the so-called ‘discoveries’ of the proliferating ‘neuro’-fields—neuroethics, neurosociology, neurohistory, neuromarketing, and beyond—seek to displace other ways of understanding human nature. They do so largely on the strength of the new “kinds of images that, since the 1990s, have flooded the public domain and rapidly acquired iconic value,” images produced by new technologies for mapping the living brain such as PET and fMRI (Vidal, 2009, pp. 24–25; Dumit, 2003; Racine, Bar-Ilan, & Illes, 2005). Out in the public sphere, these can function as ‘immutable mobiles,’ in Latour’s sense, serving to enhance the credibility of those who can marshal them effectively. Behind the scenes, however, they depend on an enormous labour of averaging and aggregating, both the firings of many neurons into ‘voxels,’ and the scans of many brains into reference ‘atlases’ for further imaging studies. Thus as Anne Beaulieu argues in a study of these atlases, “‘mobiles,’ especially digital ones, are perhaps not so much ‘immutable,’ as recreated according to conventions, and allowed to circulate in the locales where these conventions are operating” (Beaulieu, 2001, p. 655). Like the story of ‘information-

processing psychology,' and that of the informational turn in genetics (Kay, 2000), the rise of 'neuroinformatics' is a matter of the increasing deployment of computational tools in a scientific field. This development, however, has been far from an unquestioned success. Below, I discuss some of the interactions between the psychological and the neurophysiological in current research.

There has certainly been no shortage of institutional support for these initiatives: starting in 1990, then-President George H.W. Bush declared an interagency initiative known as the 'Decade of the Brain,' declaring that by "mapping the brain's biochemical circuitry," neuroscientists may help produce "more effective drugs for alleviating the suffering of those who have Alzheimer's or Parkinson's disease," while at the same time their research "may also prove valuable in our war on drugs, as studies provide greater insight into how people become addicted to drugs and how drugs affect the brain" (Bush, 1990). Though in itself merely an 'awareness' initiative, carrying no additional funding, the 'Decade of the Brain' signaled a broad prioritization of brain research by U.S. funding agencies. It was certainly successful in increasing the public "visibility of neuroscience" (E. G. Jones, 1999). Where its initiatives seemed to be most successful, however, was when they set aside the goal of understanding cognition and the mind, to focus instead on the construction of new tools for integrating and analyzing many sources of neuroimaging data.⁶⁷ Such atlases, instruments, and software suites have spurred on the explosion of interest in neurobiology and allowed for a proliferation of claims about associations between specific brain regions and behaviours of interest. Neuroscientists view this as a steady accumulation of knowledge, which despite the challenges of

⁶⁷ As one grant proposal from this period puts it, "We make no assumptions about the relationship between structure and functions in the human brain, at either a macro- or microscopic level, except to state the obvious, that these relationships are complex and poorly understood. Further, we are not proposing that we will unravel this complexity with the data collected in the context of the consortium program. The development of a probabilistic reference system and atlas for the human brain simply provides the framework in which to place these ever-accumulating data sets in a fashion that allows them to be related to one another and that begins to provide insights into the relationship between micro- and macroscopic structure and function" (in Beaulieu, 2001, p. 661).

wrangling large data sets, will eventually lead to a solid understanding of how the brain produces all the phenomena of cognition, mind, and behaviour.

Often hailing from psychology and other constituent disciplines of cognitive science, others have argued that despite their usefulness for anatomical and physiological purposes, extant neuroimaging technologies are unlikely to produce such an understanding, and that because of the layers of computational mediation required, many studies claiming to localize mental phenomena in the brain may be spurious (Uttal, 2003, 2011).⁶⁸ Critiques raised since the Decade of the Brain both by critical scholars and by cognitive scientists, cautioning against undue privileging or naturalizing of findings grounded in neuroimaging, have considerable merit. The American Psychological Association, for its part, attempted to christen the following decade the ‘Decade of Behavior;’⁶⁹ another group of researchers later called for a ‘Decade of the Mind,’ an initiative to build “new tools that can deeply probe mental processes,” particularly on “aspects of the mind believed to be uniquely human, such as the notion of self, rational thought processes, theory of mind, language, and higher order consciousness” (Albus et al., 2007, p. 1321). Of particular interest in this regard were tools for “modeling the mind,” by “combining theoretical and computational methodologies with empirical findings... for healing, understanding, and enriching the mind” (*ibid.*). Both proposals met with little official support. Instead, President Barack Obama announced another major decade-long funding program for the neurosciences in 2013, the BRAIN initiative; in the same year, the

⁶⁸ On the specific topic of attention, for instance, Uttal contends that in brain imaging studies of attentional processes “at least some investigators have reported activations in virtually all parts of the brain,” and that any greater prevalence of activations in the parietal lobe or any other location may be either a ‘culture of science’ artifact—where early discoveries “tended to direct the focus of the collegial community to that part of the brain”—or a technological one, where a particular imaging technology or calibrating atlas works better on some areas than others (Uttal, 2011, pp. 261, 263).

⁶⁹ Despite some efforts at building connections with U.S. politicians, the APA’s proposal seems to have been considerably less successful, both as an ‘awareness’ campaign and an attempt to marshal funding support. The slogan never received Presidential affirmation, as the Decade of the Brain had, and at the time of writing, the official APA website for the initiative is no longer functional (www.decadeofbehavior.org).

‘Human Brain Project’ (HBP, by analogy of course with the wildly successful Human Genome Project) was launched in the EU, a project which was eventually approved as a ‘flagship’ project slated to receive up to 1 billion Euros in funding.

Although the brain has once again been established as an essential reference point for understanding cognition, the longstanding focus of cognitive science on modeling the mind rather than measuring the brain still manifests in important and controversial ways. These new projects have also once again been the object of considerable hype, with the added twist that now many of the debates which might once have remained the subject of informal ‘corridor talk’ amongst researchers are now being aired on the World Wide Web. The remainder of this section focuses on three such projects, and some of the related press coverage: one of these is the European Blue Brain Project, at the heart of the HBP; another is the American DARPA Synapse project, led by Dharmendra Modha, a signatory of the ‘Decade of the Mind’ proposal; the last is the Canadian ‘SPAUN’ model, based on open-source and distributed-computing principles. I use these three examples to emphasize both the continuing shared focus on computers as models for the mind, and the distinct ways in which contemporary modeling efforts leverage new developments in neuroscience and computing. The Human Brain Project is the biggest and most ambitious of these; it is also the one which seeks to most closely engage with neurobiology, with its primary objective being to “develop ICT tools to generate high-fidelity digital reconstructions and simulations of the mouse brain, and ultimately the human brain” (Human Brain Project, 2012). This focus, and its polarizing leader Henry Markram, have also made it the most controversial of the three. The broader HBP is largely an outgrowth of the priorities he developed within his earlier ‘Blue Brain’ project, which in turn was fuelled by Markram’s desire to understand his son’s autism diagnosis, and a

dissatisfaction with the current state of psychiatric knowledge.⁷⁰ This compelling personal narrative, and his charismatic advocacy for the modeling programme, made Blue Brain the object of widespread media interest—echoing earlier fascination with the ‘clicking brains’ of cybernetics—and his team developed into the core of the largest collaborative neuroscience project in Europe today (H. Markram, 2009; Lehrer, 2011; Requarth, 2013).

Where most connectionist models employ at most a few thousands of interconnected nodes (often far less) and a highly simplified, abstracted ‘neuron-like’ element, the Blue Brain and HBP

projects aim to construct increasingly complete simulations of the brain, both in the number of neurons and the detailed functioning of each. Markram lays out a path of linear progress in public proclamations (Figure 22), whereby increased

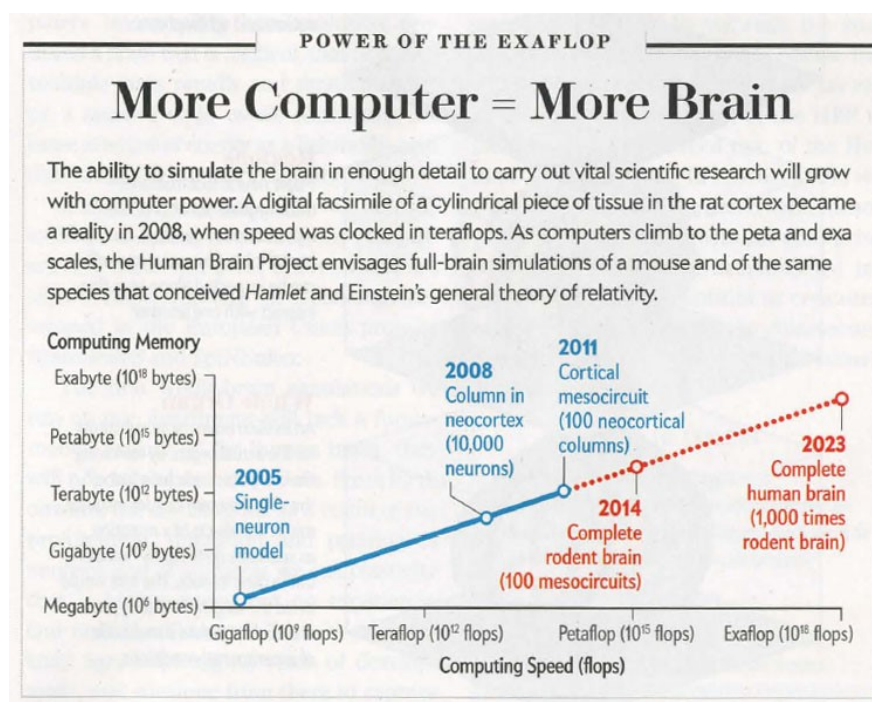


Figure 22 – Henry Markram's path of progress

⁷⁰ As one journalist recounts, “over time, trying to understand [his son] Kai's autism became his obsession... ‘We went all over the world and had him tested, and everybody had a different interpretation,’ Markram says. As a scientist who prizes rigor, this infuriated him. He'd left medical school to pursue neuroscience because he disliked psychiatry's vagueness. ‘I was very disappointed in how psychiatry operates,’ he says. It drove what he calls his ‘impatience’ to model the brain: He felt neuroscience was too piecemeal and could not progress without bringing more data together. ‘I wasn't satisfied with understanding fragments of things in the brain; we have to understand everything,’ he says. ‘Every molecule, every gene, every cell. You can't leave anything out’ (Szalavitz, 2013). The computational brain modeling is intended to link up with autism research by way of laboratory research using mice deliberately given ‘autism-like’ symptoms by prenatal exposure to valproic acid, a drug known to cause increased risk of autism when taken in pregnancy by human mothers as well. Markram and his wife Kamila are also known for their promotion of a controversial ‘intense world’ theory of autism and corresponding set of proposed interventions (K. Markram & Markram, 2010)

supercomputing power will allow his team to construct a simulation of the whole human brain by the year 2023 (H. Markram, 2012, p. 37). Their initial test case, however, is the neocortical column, best imagined as “a cylinder of tissue about half a millimeter in diameter and 1.5mm in height,” made up of six vertical layers and roughly 10,000 neurons of a few hundred different types. As he puts it, using a venerable old metaphor, the organization of their interconnections “resembles the way telephone calls are assigned a numerical address and routed through an exchange” (*ibid.*). With extensive material support from IBM, using its Blue Gene-L supercomputing system, the project has succeeded thus far in functionally replicating 100 of these columns; while its focus is primarily on modeling biology rather than producing behaviour, the columns have been ‘trained’ on certain simulated model tasks, such as keeping a ball roughly in the middle of a circular tray as it rolls

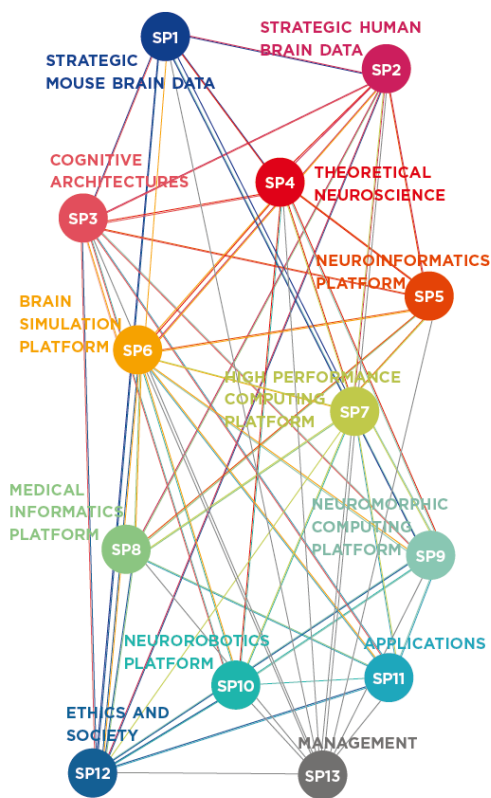


Figure 23 - Human Brain project graphic

around unpredictably. Markram’s Blue Brain group presents itself as one of more than a hundred international collaborators within the HBP, working together to follow a ‘roadmap’ that includes not only brain simulation but medical applications, ‘neurorobotics,’ and new applied computing architectures (Human Brain Project, 2012). As I discuss below, he has also been outspoken in his criticism of other approaches linking neuroscience and computer modeling. Yet some within the neuroscience community have lately begun to mobilize against Markram and the

HBP, arguing that its governance is opaque, its funding ill-directed, and the very aim of simulating the brain misguided and premature.

This controversy has been covered in detail by a growing circle of scientific journalists. The impetus was an ‘open letter’ of the type only possible online: criticizing the “narrowing of goals and funding allocation” within a second round of HBP grants, including the removal of a basic-neuroscience subproject and several associated laboratories, a group of 156 researchers posted their complaints online at ‘neurofuture.eu.’ By the time of writing, many more had signed on to the complaint after the fact, with the list expanding to 778, and coverage has continued of a potential HBP ‘boycott’ by the international research community (Sample, 2014). Part of the issue is that many neuroscientists believe our knowledge of the brain is simply too immature, and we are likely to build gross misconceptions into our models; what is needed instead is a systematic effort to map the ‘connectome’ using computer-aided microscopy of brain slices.⁷¹ Markram’s group incorporates data about neuron types (morphology) and the kinds of electrical signals they produce, but they follow “Peters’ Rule,” the premise, much as Ashby had proposed, that their initial connectivity is random (Seung, 2012). They argue that the important connections are shaped adaptively by the network’s firings; advocates of empirical ‘connectomics’ as a necessary prerequisite to simulation suggest instead that many important brain systems involve hard-wired connections. Markram contends “the problem is that we’re dealing with a cultural change,” and that it would take centuries to gather sufficient information about the full spectrum of variation in real human brains—so why not ‘impatiently’ proceed to simulation? As he puts it, “Maybe we see the value in their data and they

⁷¹ One of the most interesting aspects of the Blue Brain project is the presence of filmmaker Noah Hutton, who is working on a documentary about the group, and releasing it in stages; it equally covers critics of the group, such as Sebastian Seung (Hutton, 2012), at present the foremost popularizer of the connectome concept and advocate of the microscopic analysis approach (Seung, 2012, 2013)

don't see the value in our models, you know, okay! It takes time... the Blue Brain Project doesn't compete with anybody, there is no competition. I don't see what anybody is doing out there as a competition" (in Hutton, 2012). Markram invokes an *ethos* of scientific collaboration and platform-building in admirable and persuasive ways.

In practical terms, however, Blue Brain has been a part of several grant competitions, and its brain-simulation program stands at the core of the Human Brain Project, which through such a competition became one of only two 'Future and Emerging Technologies' projects to receive the highest 'flagship' level of funding from the European Commission, as part of its "Digital Agenda for Europe" (European Commission, 2014).⁷² Entwined with but distinct from the intellectual debates over the state of neurobiological knowledge, at issue in the present controversy is the process by which the Human Brain Project doles out over a billion Euros in funding, and precisely whether this is sufficiently competitive. As one signatory to the petition put it, the problems were first of all "a huge amount of salesmanship, which the HBP now admits was 'miscommunication' about the true goals and nature of the project... that leaves us with a €1.2 billion [*sic*] for a project that has very modest and unclear deliverables and is not going to reach the lofty goals it promised. It leaves the neuroscientists who are in the project without a clear goal other than 'providing data.'" Compounding this issue, he argues, is a "very problematic system for funding and managing science," where the "executive committee make the decisions on who receives funding, including how much money goes to themselves" (Neuroskeptic, 2014).

The petition signatories have pointed to the American BRAIN initiative as embodying a better approach, with more transparency as well as better separation between those who led the grant application, and those who will eventually dole out its funding. This is intertwined with the

⁷² The other, incidentally, deals with nanomaterials, and seems to have produced far less public controversy.

views of many scientists internationally about broader problems with the “mega-project model” and the “broken publishing and promotion systems” (*ibid.*). In this sense, it is at once a debate about how to properly conduct the ‘internal’ processes of scientific research, in terms of what techniques and materials one employs, and about the proper conduct of its ‘external’ administrative, institutional processes, in terms of how funding should be granted. The two domains are distinct, but in practice difficult to disentangle. All of the researchers I discuss here are united by their belief in the application of computational tools to understanding the brain and mind, but at the same time they are engaged in ongoing and multifaceted debates about how best to do so.

Markram advocates detailed simulations of neuronal columns, incorporating as much biological realism as possible. He persuasively argues that this will allow us to move beyond the vagaries of imaging particular brains and capturing their variations, as well as allowing us to ‘zoom in’ to the level of individual neurons, unlike the rather broad measurements captured in the voxels of functional imaging. Other neuroscientists contend that we just don’t know enough yet about the biological brain to proceed to this simulation stage, and instead we need more functional imaging, more microscopy of brain slides, and more use of ‘big data’ and neuroinformatics techniques to synthesize these data sets.⁷³ Another site, however, has also figured prominently in public controversies over brain modelling: the ‘cognitive computing’ initiative led by Dharmendra Modha at IBM’s Almaden research centre, with support from DARPA. Rather than maximal biological detail, Modha’s team takes as its primary aim the construction of new computer chip designs for practical applications, based not on the same von Neumann architecture as the majority of current digital computers, but a neuron-like parallel design. ‘Cognitive computing’ and ‘neuromorphic

⁷³ One of the most fascinating approaches they advocate is the ‘gamification’ of neuroscience data, using an online game called ‘Eyewire’ to help trace out connections between neurons. I return to this novel form of ‘citizen science’ in the next chapter.

computing’ are synonyms for this new approach to designing computers by analogy with biology; another name sometimes adopted by the field, ‘biometaphorical computing,’ seems perhaps most apt of all. Far from constituting an entirely new paradigm, however, this is really just cybernetic modelling revitalized by contemporary tools, and marks a return to an area of research which fascinated von Neumann himself (Von Neumann, 2012). Blue Brain, conversely, was designed around IBM’s Blue Gene/L supercomputing architecture—as the name suggests, it originally designed for use in bioinformatics, building on an earlier design for quantum chromodynamics research—which though massively parallel is still based on PowerPC chips with a sequential von Neumann architecture. In other words, whereas Blue Brain simulates neural networks in a program running on thousands of interconnected digital computers, Modha’s project is part of an effort to build a different kind of supercomputer around ‘neurosynaptic cores,’ making it parallel and neuron-like ‘all the way down,’ so to speak.

The trade-off is that they use a very simple ‘leaky integrate-and-fire’ model for each neuronal element, really just a short mathematical equation (Artificial Brains, 2013). Like Markram, Modha is something of a salesman, and likes to describe his project in interviews as seeking to “understand the mind by reverse-engineering the brain” (Orca, 2009). His group is the primary site for the DARPA SyNAPSE program, which seeks to coordinate “aggressive technology development activities in hardware, architecture and simulation” (Defense Sciences Office, 2013), with the ambitious goal of building a system that “matches a mammalian brain in function, size, and power consumption,” with 10 billion neurons and 100 trillion synapses, while occupying less than two liters of space and consuming less than one kilowatt of energy (Artificial Brains, 2013).⁷⁴ Modha’s project has garnered

⁷⁴ ‘SyNAPSE’ is one of the military’s beloved backronyms, standing for “Systems of Neuromorphic Adaptive Plastic Scalable Electronics.” It should also be noted that the ‘cat-scale’ simulation cited above in fact also ran on the Blue Gene architecture, not on any of their ‘neuromorphic’ chip designs, which have yet to be scaled up to that degree. The Blue

significant public interest, particularly with his team's 2009 announcement of having produced a simulation that matched the scale of a cat brain. Unsurprisingly, this led to a dispute with Markram, whose group was attempting to model just a portion of a mouse brain in painstaking detail. Just as unsurprisingly, the led to many quips about 'cat fights' by scientific journalists (Adee, 2009). What *is* perhaps surprising, due to the proliferation of 'niche' audiences within the current media landscape, was the volume of coverage, as this highly technical controversy played out in very public fashion. After Modha's group received the Gordon Bell Prize for their efforts, Markram fired off a letter to the Chief Technology Officer of IBM, Bernard Meyerson, sending copies to assorted media outlets as well. The letter constitutes a fascinating attempt to recruit—and be widely seen recruiting—the primary provider of material support for both teams:

Dear Bernie,

You told me you would string this guy up by the toes the last time Mohda [sic] made his stupid statement about simulating the mouse's brain.

I thought that having gone through Blue Brain so carefully, journalists would be able to recognize that what IBM reported is a scam ... I am absolutely shocked at this announcement. Not because it is any kind of technical feat, but because of the mass deception of the public.

These are point neurons (missing 99.999% of the brain; no branches; no detailed ion channels; the simplest possible equation you can imagine to simulate a neuron, totally trivial synapses; and using the STDP learning rule I discovered in this way is also a joke)... It is really no big deal to simulate a billion points interacting if you have a big enough computer. The only step here is that they have at their disposal a big computer. For a grown up 'researcher' to get excited because one can simulate billions of points interacting is ludicrous... I suppose it is up to me to let the 'cat out of the bag' about this outright deception of the public. Competition is great, but this is a disgrace and extremely harmful to the field. Obviously Mohda would like to claim he simulated the Human brain next - I really hope someone does some scientific and ethical checking up on this guy (in Adee, 2009).

This excerpt indicates not only the central points at issue—all surrounding the biological plausibility of Modha's neurons—but the sharp contrast between the diplomatic version of the Blue Brain project, and the polemical version which appears when institutional support for disparate

Gene supercomputer, by contrast with the SyNAPSE goals, consumes vastly more energy and in most configurations occupies a fairly large room. The ideal, however, is that building a system around neurosynaptic cores would make it better-optimized for parallel processing, hence more resource-efficient.

approaches is at stake. The debate indicates not only the multiple and ongoing points of collaboration between science, industry and the military, but also the internal heterogeneity of viewpoints within these institutions. Advocates of connectomics and other approaches more focused on measuring than modeling contend for their part that both Blue Brain and SyNAPSE have substantial problems when evaluated in terms of biological realism.⁷⁵

The open letter marked a “new low point in Markram’s relationship with IBM:” they had “started out allies in 2005, when IBM signed an agreement with Markram’s institution, the École Polytechnique Fédérale in Lausanne, Switzerland” to showcase their Blue Gene system with a brain simulation (Seung, 2012). But the rise of the cognitive computing project at Almaden led to a ‘souring’ of that relationship. As an in-house research project, and the more clearly application-oriented of the two, it seems at present that Modha’s group is the more favoured, and Meyerson never responded publicly to Markram’s pointed critiques. Though his indignation at the credulously-reported claims of having ‘simulated a cat brain’ is understandable, Markram seems to ignore the different potential aims of simulation, and the differing levels of realism needed to achieve them. The name of ‘cognitive computing’ is as much a branding effort as a scientific subfield, serving to link together projects like Modha’s with the signal achievement of contemporary commercial AI research, IBM’s Watson. Just as Deep Blue had beaten Kasparov at chess, in 2011 Watson defeated the two most successful past human contestants on the trivia game show *Jeopardy!* This project aimed

⁷⁵ Specifically, as Seung argues, we lack sufficient information even to produce a fully accurate simulation of the nervous system in a simple model organism like the worm *C. elegans*: while Markram has modeled the electrical properties of many neuron types using ‘compartments’ that simulate the ion channels of real neurons, as discussed above he incorporates no information about their interconnections, using a random initial configuration. And while both Markram and Modha have included “reweighting using mathematical models of Hebbian synaptic plasticity,” they lack models of neural “reconnection, rewiring, and regeneration” (Seung, 2012).

more for psychological than biological plausibility, and by some well-informed accounts achieved that goal.⁷⁶

Far from implementing some radically new algorithms, however, it simply put many techniques for language-processing and database-retrieval together effectively: resolving the ‘combinatorial explosion’ by not only massively parallel, extremely powerful computer hardware, but by a massive distributed human labour of describing and categorizing, in the form of the World Wide Web, and especially Wikipedia. After its televised victory, IBM announced the formation of its cognitive computing business unit around these achievements, pitching their services to the corporate world with a panoply of buzzwords: ‘artificial intelligence’ and ‘collective intelligence’ meets ‘business intelligence.’ By supporting a wide spectrum of research in this area, IBM hopes to position itself as an obligatory passage point for a new era of ‘human-computer symbiosis,’ and its public relations material aims to assuage the concerns which led it to abandon AI work in the 1950s.⁷⁷ This sense of cognitive computing introduces a new criterion of success. If Watson should prove useful in diagnosing and treating medical conditions, then does it matter whether it’s really a ‘thinking machine,’ or simply condensed and automated human expertise? If DARPA SyNAPSE should employ biological metaphors to produce a new computer architecture more effective for parallel processing tasks, then does it matter if its elements are “really” like neurons?

⁷⁶ As Watson’s opponent and former *Jeopardy!* champion Ken Jennings put it, “I expected Watson’s bag of cognitive tricks to be fairly shallow, but I felt an uneasy sense of familiarity as its programmers briefed us before the big match: The computer’s techniques for unraveling *Jeopardy!* clues sounded just like mine. That machine zeroes in on key words in a clue, then combs its memory (in Watson’s case, a 15-terabyte data bank of human knowledge) for clusters of associations with those words. It rigorously checks the top hits against all the contextual information it can muster: the category name; the kind of answer being sought; the time, place, and gender hinted at in the clue; and so on. And when it feels “sure” enough, it decides to buzz. This is all an instant, intuitive process for a human *Jeopardy!* player, but I felt convinced that under the hood my brain was doing more or less the same thing” (Jennings, 2011).

⁷⁷ One example: “Far from replacing our thinking, cognitive systems will extend our cognition and free us to think more creatively. In so doing, they will speed innovations and ultimately help build a Smarter Planet” (IBM Research, 2013).




Figure 24 - Promotional graphic for 'neuromorphic' computing, IBM Research

While disputes over funding and the proper scope of claims made in the press are inevitable, the more collaborative view of science is perhaps closer to the truth, in addition to standing as a useful regulative ideal. From the standpoint of connectomics advocates, Blue Brain captures a somewhat better but still inadequate amount of detail with its model; from the standpoint of Blue Brain advocates, Modha's chips should not even be mentioned in the same sentence as brains; from Modha's standpoint, his team is building the architecture for a new computing paradigm, and biological detail is secondary to applied usefulness. Both Markram and Modha's approaches seem to have considerable merit on their own terms, as do approaches more focused on studying the brain, not to mention the traditional research priorities of psychology. Equally, as IBM themselves contend, there are some applications where traditional, digital computers are best suited, and others where neuromorphic, analog computing is preferable—the real challenge is deciding between the two, and building them together in hybrid systems.

One final approach worth highlighting is that of the 'Semantic Pointer Architecture Unified Network' (SPAUN) group at the University of Waterloo, which has likewise received considerable popular press coverage for a large-scale brain modeling project. This is distinct from Blue Brain and SyNAPSE in two crucial ways: first, it is an open-source, freely available project intended to run on networks of widely available commodity PCs, as opposed to specialized supercomputing

architectures; and second, rather than aiming to simulate small neural systems in detail, it is “centrally directed to bridging the brain-behavior gap,” linking together 2.5 million spiking neurons to produce “a wide variety of behaviorally relevant functions,” such as Raven’s matrices, a common psychological test (Eliasmith et al., 2012). As the SPAUN group framed the difference in an ‘Ask Me Anything’ session on the popular online forum Reddit, Blue Brain aims to

simulate, as realistically as possible, the number of neurons in a human brain. What we’re more concerned with here is how to hook up those neurons to each other such that we get interesting function out of our models, so we’re very concerned with the overall system architecture and structure. And that’s how we can get out these really neat results with only 2.5 million neurons (which is just a fraction of the 10 billion a human brain has). We are definitely interested in scaling up the number of neurons we can simulate, but it’s secondary to producing function. (DeWolf & Stewart, 2012)

The prioritization of psychological functioning and task-trainability in their model goes along with an emphasis on producing standardized, freely shareable research platforms for cognitive science: the ‘Semantic Pointer Architecture’ and the ‘Neural Engineering Framework’ (Eliasmith & Trujillo, 2014, p. 2). SPAUN is currently the largest implementation of these platforms, running on a Canadian distributed computing system called SHARC. While far from a complete simulation of the human brain, it uses neuroscientific data to model a complete pathway from optical perception, through various processing steps, to motor output through a simulated ‘arm’ it uses to write.

Chris Eliasmith of the SPAUN group has also been a public advocate of this particular approach to modeling, with a quasi-popular book entitled ‘How to Build a Brain,’ which serves equally as a practical introduction to coding within the open-source Neural Engineering Framework (Eliasmith, 2013). He has also presented a persuasive analysis of the debates between Markram, Modha, and other proponents and detractors of large-scale brain modeling. Eliasmith argues that without such models we cannot properly address the functioning of the brain or form “large-scale

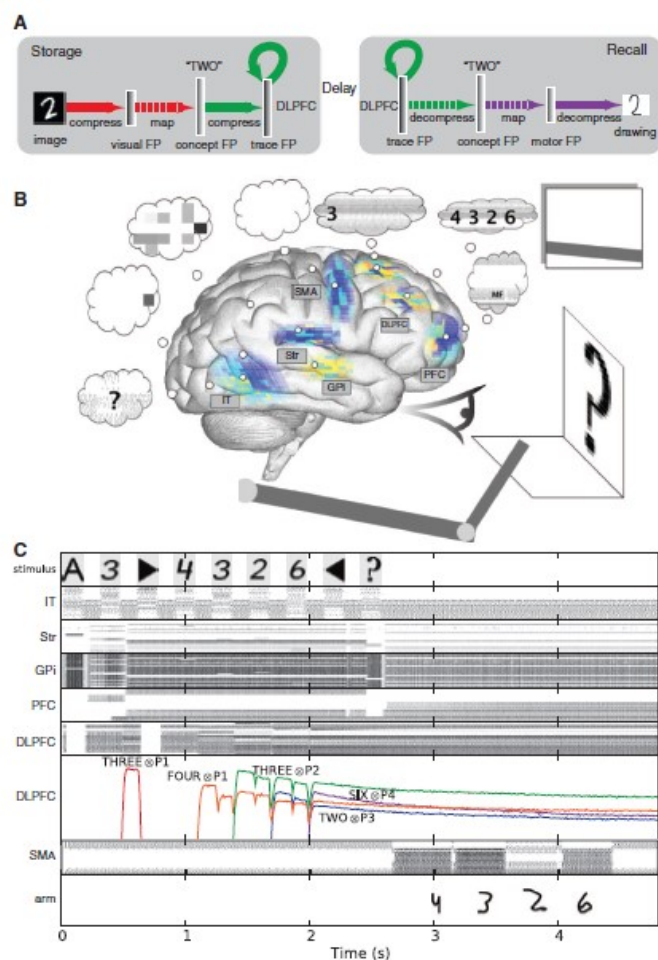


Figure 25 - SPAUN and the 'Serial WM' test

hypotheses” about how the mind emerges from the brain, and notes that the fundamental issue in these debates is over the “right ‘level of detail’” to be incorporated in such models. This is an “ill-posed question,” he contends, and in fact “as has long been accepted by those constructing large-scale climate models, the appropriate scale is determined by balancing two things: first, the questions that need to be answered and second, the available computational resources” (Eliasmith & Trujillo, 2014, p. 3). There is value, they recognize, in both bottom-up and top-down approaches; the former, as

epitomized by the Human Brain Project, seeks to simulate biological process in full detail, while the latter as adopted by SPAUN “entails identifying hypotheses regarding the behavioral function of a brain area and then determining how neurons carry out the relevant computations with networks of spiking neurons” (*ibid.*, p.4). The bottom-up approach is only as good as the neuroscientific data one feeds in, making it susceptible to critiques as outlined above. The top-down approach can be biased by assumptions and pre-existing hypotheses, but it can be tested against the “vast knowledge gained through the behavioural sciences” respecting human behaviour and task performance. We do not know in advance the right level of detail for modeling the brain, but “exploring plausible levels in

the context of behavior is likely to lead, most efficiently, to a good understanding of the structure/function relation” (*ibid.*).

In other words, if your model of 2.5 million simple spiking neurons can perform a task like Raven’s matrices in a way that matches up with observations of humans doing the same thing, then you can be confident you have reached a sufficient level of detail. If the behaviour of your model varies greatly from what you observe in humans, you’ve likely missed something. Success or failure in cognitive tasks, on this view, is the best measure for a model of cognition. By contrast, when you are simply synthesizing basic data about neurobiology into a computer program, the criteria for success—as the signatories of the petition against the HBP lament—become rather vague.⁷⁸ As is often the case in bitter intellectual controversies, these three projects have more in common than their differences would suggest. All aim toward simulating the brain and ultimately the mind using computational tools; the differences among them are primarily with respect to scales and levels of description. What in my view associates them most closely with the historical development of cognitive science is their prioritization of computer models as objects of inquiry in their own right, rather than as mere means of accessing the biological brain or the psychological subject. This is at the heart of my argument that cognition as articulated by cognitive science is a profoundly *mediated* cognition. Its corollary is that this research program has long garnered considerable interest in the popular media, a phenomenon which only seems to be increasing. These tool-oriented programs of simulation are indeed controversial, especially among those advocating for more traditional disciplinary approaches. But they may offer a way to move beyond the localizationist and essentialist

⁷⁸ As discussed above, the Blue Brain project does at times train their networks to achieve simple tasks in simulated environments, like maintaining the position of a rolling ball in the centre of a tray. These may give some insight into general dynamics of neural firing, but cannot be plausibly conceived as analogous to human completion of the same task: their program uses four simulated ‘muscles’ attached to the corners of the platform, and evidently human completion of the task involves distributed function across the brain, not isolated to a single small ‘column.’

tendencies sometimes present in popular and professional discourse of ‘brainhood:’ revealing a human nature that, far from being fixed at birth and delimited according to specific brain areas, is highly malleable and broadly distributed. Rather than seeking the one true approach to modeling cognition, then, and pursuing it to the exclusion of all others, we should be open to a similar sort of plasticity in the pathways of scientific inquiry. Specific findings with regard to neuroplasticity, however, are of course still up for debate (Eklund, Nichols, & Knutsson, 2016).

No discussion of where we stand today with respect to artificial cognition and ‘human-computer symbiosis’ would be complete without mention of Google. It is not merely the ‘search engine’ which has come to define our era of the Internet, but the truest realization of both dreams thus far. I have interacted far more with Google in the course of writing this text than with any human, and it has been an integral component of the cybernetic system which constructed this argument. It cannot converse with me, exactly, and certainly not in any way that would pass the standard set by Turing’s imitation game—although its voice recognition and synthesis on mobile operating systems are increasingly reliable. Far more usefully, however, with its help I can find just about anything on the basis of very limited information. In one of many similar episodes while composing this chapter, I recently typed ‘von Neumann paper component reliability,’ and was returned the paper in question as my second result (Von Neumann, 1956), with another article commenting on the same as my first. Querying ‘movie about sled’ returns ‘Citizen Kane;’ ‘song about girl with birthmarks’ gives the correct name and accompanying video for the tune.⁷⁹ I can ask how long it will take me to get to work, and it will give me an estimate and live traffic information for a drive to campus, based on its stored profile for me. If it knows the traffic is particularly heavy it will, unbidden, pop up a notification on my phone and my computer simultaneously. In these and

⁷⁹ 1993 hit “Mmm Mmm Mmm Mmm” by the Crash Test Dummies.

so many other ways, Google has become the most prominent cognitive *prosthesis*—to paraphrase McLuhan—for many of us in the developed world. But it is definitively not a single ‘artificially intelligent’ software program. Rather, it is an assemblage of distributed and networked algorithms, aggregating data produced and categorized by the multitudes of human Web users. IBM’s Watson, as I described above, operates in much the same way. The real genius of Google’s PageRank algorithm, what made it such an enormous improvement over earlier search engines, was that instead of trying to collect and analyze increasingly fine-grained detail about the content of an ever-increasing pool of Web sites, it would gather information about the *connections between them*. The patterns of linking between pages could be distilled into a representation of what people have interpreted them as being about. Thus von Neumann’s paper was returned as a top result because many others have linked to that document on their own pages with nearby text referring to component reliability.

George Dyson has interpreted the success of Google as an indication that “in the age of all things digital we are building analog computers again,” a line of research which Turing and von Neumann, “for all their contributions to the digital revolution, did not see...as a dead end.” Despite being implemented on massively parallel, internetworked digital machines, “it is in the analog domain that the interesting computation is being performed” by search engines and social networks, with information being “encoded (and operated upon) as continuous (and noise-tolerant) variables such as frequencies (of connection or occurrence) and the topology of what connects where... Pulse-frequency coding for the Internet is one way to describe the working architecture of a search engine, and PageRank for neurons is one way to describe the working architecture of the brain.” Whether implemented in a biological system or an electronic one, these both constitute distinctive kinds of information-processing structures, reliant on digital components but as a system

performing analog computation, using continuously variable software elements (G. Dyson, 2012). These intriguing reflections highlight the enduring appeal of the mind-computation analogy, while suggesting some of the ways its parameters have begun to shift, given new technological reference points and increasingly refined hybridizations of the digital and the analog.

Wired editor Kevin Kelly describes systems like Watson and Google as fully realized ‘AI without consciousness.’ He reports that Google’s cofounder Larry Page stated quite matter-of-factly from the company’s origins that they were “really making an AI” – a plan that’s taken a further step recently with the restructuring of the company as ‘Alphabet,’ a holding company whose major profit centre is Google search and advertising, but which operates separate divisions with their own CEOs pursuing more speculative ventures like self-driving cars, home automation, medical technologies, and life extension. All of these involve elements of AI, and all aim to further solidify Google’s market-leading status as provider of cognitive prostheses. As Kelly believes,

rather than use AI to make its search better, Google is using search to make its AI better. Every time you type a query, click on a search-generated link, or create a link on the web, you are training the Google AI. When you type “Easter Bunny” into the image search bar and then click on the most Easter Bunny-looking image, you are teaching the AI what an Easter bunny looks like. Each of the 12.1 billion queries that Google’s 1.2 billion searchers conduct each day tutor the deep-learning AI over and over again. (Kelly, 2014)

This implies a vision of AI which is less a “discrete machine animated by a charismatic...humanlike consciousness,” as long envisioned by science fiction, and instead something “more like Amazon Web Services—cheap, reliable, industrial-grade digital smartness running behind everything” (*ibid.*). Behind this industrial smartness, of course, is cheap and reliable computing hardware, networked together in massively parallel systems. Underpinning this shift of late has been a shift from general-purpose ‘CPU’-type processors to the more specialized GPU chips developed for processing graphics, which are better adapted to carrying out many parallel calculations simultaneously. Networked together in clusters, GPU computing is useful for simulation in physics—in effect,

modern computer graphics processing involves simulating the paths of millions of photons—but also myriad other fields and commercial applications. Bioinformatics and protein modelling suites are being adapted to GPU operation with enormous performance gains. The SPAUN architecture for cognitive modelling is also designed to run on a networked GPU system, and this can in fact be purchased as a ‘cloud service’ (Eliasmith *et al.*, 2012). Equally, “neural nets running on GPUs are routinely used by cloud-enabled companies such as Facebook to identify your friends in photos or, in the case of Netflix, to make reliable recommendations” (Kelly, 2014) for films you might like to watch. The work of Modha’s group, seen in this light, is really seeking to eventually supplant GPU computing, making still more parallel and now purposefully neuron-like computing hardware into an equally commonplace, cheap, and reliable component slotted into the future infrastructure of cognition-as-software-service.

Debates will continue to rage between advocates of connectionism and old-fashioned cognitivism, between the more biologically-oriented and the more abstraction-oriented, between those who privilege mind or brain, analog or digital. Their resolution is less a matter of universally-applicable discovery than of particular decisions, of which classes of system different researchers choose to adopt, and what they can accomplish with them. Any resolution to the controversies, if ever, will come only with the construction of a truly intelligent *and* conscious system – if consensus could even be achieved as to when that goal were fulfilled. What has come the closest to achieving artificial intelligence for now are systems like Google and Watson, non-conscious hybrids of digital and analog computing, of distributed, networked computing hardware and human expertise. This is human-computer symbiosis in its most highly-developed form each day; we are collectively advancing the experiment with every query we make and everything we make ourselves in dialogue with the computer. Cognitive scientists, as we saw in Licklider’s writings and elsewhere, have often

been highly reflexive about the ways that technologies worked to shape their thinking and research practices. As for what scientific research has been done investigating the broader question of media effects, and how in turn that has been covered in the popular media, I take these questions up in the following sections.

Summing up, then, my goals across this broad-ranging and eclectic history have been to foreground tools and technological mediation as a way of understanding the pluralistic, ‘transdisciplinary’ field of cognitive science, and to offer my own taxonomy of three distinct eras within this history. Changes in theory have always been closely coupled with changes in the hardware and software of computing tools, from the analog control systems of the cybernetic era, to the elevation of digital computing by cognitivism, and the multiplying hybrids of today. ‘Progress’ in the field has meant different things to different actors, and has always involved reciprocal translations between not only technology and theory, but practice and metaphor, materiality and abstraction. I have tried to grant equal emphasis to the many elements of heterogeneity across the history, and to an underlying unity as well. Delineating cognitive science proper, as opposed to the various related fields of neuroscience, psychology, linguistics, or AI, is a focus on understanding the mind, and a conviction that the best way to do so is through the mediation of computing. Computers are thought, moreover, to offer access directly to the phenomena of mind and cognition, rather than serving as intermediaries for the study of brains. Controversies rage amongst the core set of researchers who continue to adopt this approach, but across its pluralistic history this root metaphor of mind-as-computer is what holds together the field of cognitive science.

4

Media effects in science and society

“All of man's artefacts - whether language, or laws, or ideas and hypotheses, or tools, or clothing, or computers - are extensions of the physical human body or the mind. Man the tool-making animal has long been engaged in extending one or another of his sense organs in such a manner as to disturb all of his other senses and faculties. But having made these experiments, men have consistently omitted to follow them with observations.”

— Marshall and Eric McLuhan (1999, p. 93)

Having considered the long history of technological mediation in cognitive science, this chapter turns more directly to the question of how media act upon the mind and brain. What scientific research has examined this question, both generally and under the specific auspices of cognitive science? It is a widely-held view that new technologies are affecting the way we think, as I found in most informal discussions of this dissertation project, along with several recent popular books (eg. Brockman, 2011; Carr, 2011) and myriad journalistic articles. According to a Pew survey of American educators, 87% of those surveyed believed that “widespread Internet use was creating an ‘easily distracted generation with short attention spans,’” and 88% that “today’s students have fundamentally different cognitive skills because of the digital technologies they have grown up with” (Mills, 2014, p. 385).

This age-old concern for media effects is increasingly being dressed up in the new clothing of neuroscience. Nicholas Carr’s book on “what the Internet is doing to our brains” opens by citing McLuhan, and often echoes his proclamations, arguing that “media work their magic, or their mischief, on the nervous system itself. Our focus on a medium’s content can blind us to these deep

effects. We're too busy being dazzled or disturbed by the programming to notice what's going on in our heads" (Carr, 2011). Carr is one recent advocate for this position, but there is perhaps no figure more strongly associated with this view, or more polarizing, than Susan Greenfield, the British peer who has carved out a niche in the tabloid press stoking concern about the effects media on our brains. Her fears relate to the Internet in general, but also nearly any specific instance of new, particularly screen-based media—including not only Facebook, for instance, but the short-lived user interface for the Android mobile operating system developed by the social networking corporation ("Baroness Susan Greenfield," 2010; S. Greenfield, 2013). She has argued in a recent book that 'Mind Change' driven by media technologies should be ranked with climate change as among the foremost challenges facing contemporary society (S. Greenfield, 2014).

Greenfield has also been denounced as a formulaic fearmonger with little solid evidence to back up her claims. One critic offered a 'how-to guide' for disparaging any given technology in the style of Greenfield, "as long as it's something new that people can experience in a way that involves the brain." Simply point out that the brain is "very adaptive," that it may change in response to social networks, video games, or whatever, that children may potentially be harmed by these changes, and generally imply that too strong an insistence on evidence is "for bitter people who hate children" (Burnett, 2013).¹ Such critics rightly insist that Greenfield's opinion pieces overstate the rather more nuanced and limited evidence for technology changing the brain. What does this evidence look like, though, and how is it produced? In the next section, I conduct a media analysis focused directly on one specific claim: that media cause mental pathologies like ADHD. As already

¹ The author heaps further scorn upon the usual venues of publication for her pieces, in places like the *Daily Mail* and *Telegraph*, saying that you get "extra Greenfield irony points if you contextualise your claims amid wider concerns about technology making children too insular and disengaged from the outside world," then publish in a paper "that regularly portrays the outside world as a lawless maelstrom of perverts, workshy criminals and powerful carcinogens" (Burnett, 2013).

discussed in the introduction, and as further evidenced by the ‘Greenfield formula,’ this concern for harm in childhood development is a primary driver of public concern with respect to media effects. In this section, however, I examine the relevant scientific research and try to separate out the question of how media technologies may be working more fundamentally to shape the ‘normal’ mind and brain as well.

I develop a three-part argument in analysing the controversies around this question. First, I examine the history of media effects research, with particular reference to a now-canonical disciplinary history: from a ‘hypodermic needle’ model of powerful, persuasive media, to a ‘limited effects’ model which emphasized the overriding influence of small groups and ‘primary’ social ties (Pooley, 2006a, 2006b). I then consider where McLuhan and ‘enculturation’ research stand with respect to this history, and how the professional and popular discourse on media effects has shifted with the rise of the Internet and of the cognitive sciences. Two interconnected sets of concepts outlined in the previous section have become obligatory passage points for any scientific discussion of media effects: those of information-processing and of neurobiology. Arguments that new media are changing the way we think—and especially how our children think—typically invoke these scientific concepts, but in fact there is surprisingly little research to back them up. As one analysis of journalistic engagement with effects of digital media on the ‘neurological adolescent’ contends, “while the evidential basis is thin and ambiguous, it has immense social influence” (Choudhury & McKinney, 2013, p. 193). In one sense, then, part of my story here is that of a hyped-up moral panic, overselling science to sell newspapers (or, increasingly, online advertisements). My first section traces out how we came to this point, from hypodermic needles to limited effects to subconscious cultivation. The second part examines the contemporary turn toward a somewhat overblown conception of media effects as occurring on a neurobiological level. The third section

and final part of my argument, however, highlights an important causal factor in the apparent lack of compelling neurocognitive research into media effects: that technological mediation is, somewhat paradoxically, ubiquitous in these research practices.

Almost invariably, when a brain imaging study claims to be investigating some facet of human experience, it is in fact investigating some *mediated presentation* of an experience to a subject confined within a scanning apparatus. I consider some critiques and potential advantages of this approach, in the broader context of STS engagements with the construction and use of brain images. Finally, this leads into a problematization of the boundary between normal and pathological, and the conclusion that these ultimately *cannot* be fully disentangled from one another. The elusive definition of a ‘normal’ brain/mind produces difficulties in understanding how media technologies may be altering our thought processes. As others have argued, and McLuhan himself seemed to recognize, the idea of the “extended, embedded, and enacted mind can offer a decentering alternative to the current debates” (Choudhury & McKinney, 2013, p. 209), allowing us to recognize that the mind is not simply changed, for better or worse, by the use of new media. Instead the mind is fundamentally structured by its interactions with media new and old, beginning with symbolic natural language and ending with the all-encompassing hybrid form of networked digital computing. I conclude by noting that such debates may be pushed aside by a new set of research priorities in cognitive science, centred on technologies for simultaneously *intervening in* and representing the ‘rewired’ brain. While I aim to complicate and question the current discourse of media effects, the sum of my tripartite argument is that our brains and minds are indeed being affected in significant, perhaps unexpected ways by our interactions with new media.

Delimiting 'effects.'

Like many trained in media and communications studies programs, I first encountered the history of media effects research in the following form. Once upon a time, in the years surrounding the First World War, scholars of media—and particularly of *propaganda*—believed in a model of powerful media effects known as the ‘hypodermic needle’ or ‘magic bullet:’ skilled designers of persuasive messages could create content which would be accepted with minimal critical reflection by a majority of the populace, injected directly into their minds, so to speak. The famous Orson Welles broadcast of *The War of the Worlds* in 1938 is often cited as support for this view of the credulous audience, and as an impetus for a counter-movement in communications research. This was led by Paul Lazarsfeld, Herta Herzog, Elihu Katz, and others who argued for a more sharply limited or ‘two-step’ model of mediated communication, wherein the content of media was only of secondary importance in persuading public opinion, as compared with the importance of small-group ties (one’s circle of friends, family, religious or military associates, and so on.). Now enshrined in various textbooks (e.g. Croteau & Hoynes, 2006), this Whiggish history was first proposed by Lazarsfeld and Katz along with Edward Shils, and it has been most thoroughly contextualized and revised by Jefferson Pooley (Pooley, 2006a, 2006b). This internalist dialectic is often narrated with a third stage as well, in which the rise of cultural studies and George Gerbner’s cultivation model (Gerbner, 1998) marks a synthesis of sorts. As with the origins of cognitive science discussed in the previous section, the origins for this stream of communications research lie in war and military priorities: above all, to understand the effects of propaganda, ideally finding means of inoculating friendly populations against hostile campaigns, and devising more effective ways to propagandize enemy populations.

One of the first important empirical papers to promote a ‘limited effects’ model, and emphasize the importance of ‘primary groups’ – those characterized by “intimate face-to-face association and cooperation” (Cooley in Shils, 1948, p. 283) – was Shils and Morris Janowitz’s analysis of “Cohesion and Disintegration in the Wehrmacht in World War II.” They contended that

at the beginning of the second world war, many publicists and specialists in propaganda attributed almost supreme importance to psychological warfare operations. The legendary successes of Allied propaganda against the German Army at the end of the first world war and the tremendous expansion of the advertising and mass communications industries in the ensuing two decades had convinced many people that human behavior could be extensively manipulated by mass communications... [but] studies of the German Army’s morale and fighting effectiveness made during the last three years of the war throw considerable doubt on these hypotheses (Shils & Janowitz, 1948, p. 314).

Instead, despite the existence of a ‘hard core’ of politically-motivated volunteers, in this case driven by their National Socialist convictions, whether a given group of soldiers in a conscripted army would desert, ‘passively surrender,’ or fight to the death was for the most part a matter of primary group relations. If they felt a tight bond with the men of their company, and those who expressed doubts about the war or belief in Allied propaganda were ostracized as ‘bad comrades,’ soldiers were likely to fight harder and resist persuasive messages from the enemy. The point was of course not that propaganda was useless, but that its effects were limited: the “erroneous views concerning the omnipotence of propaganda must be given up and their place must be taken by much more differentiated views as to the possibilities of certain kinds of propaganda under different sets of conditions” (Shils & Janowitz, 1948, p. 315). Determining those differential possibilities was to be the task of postwar mass-communications research. Channeled through a heterogeneous set of motives, the powerful-to-limited-effects narrative made its way from the work of Shils & Janowitz to Lazarsfeld and Katz’s “simple, direct, and meagerly sourced” historical narrative in *Personal Influence*,

on to the textbooks and lecture halls of North American communications studies courses (Pooley, 2006b, p. 132).²

Lazarsfeld and Katz thus typified the stream of public opinion-oriented media theory which held that the effects of media were minimal in terms of persuading viewers on specific issues: whether to vote for this or that candidate, to buy war bonds or not, to desert, surrender, or fight to the death. I do not extensively consider this view, already well examined elsewhere, except as a critical foil to the subsequent efforts to understand how media may in fact exert significant effects on our thinking. These are best exemplified by the cultivation research of George Gerbner and the Cultural Indicators project, which focused chiefly on the *content* of media, and the work of Marshall McLuhan, which as sought to foreground the *formal* aspects of media. I argue this is one of the most important distinctions to be made amongst the attempts to understand media effects: they are not mutually exclusive, by any means, and the effects of form and content are difficult to separate, but often differences of opinion about media effects are grounded in different notions of where to look. Katz and Lazarsfeld considered the effects of specific persuasive messages, particularly as in propaganda and public-relations efforts, making them paradigmatic of an ‘administrative research’ that—much like Herbert Simon’s ‘organizational research’—sought to support the effective functioning of commercial and governmental institutions. Gerbner’s research programme looked to understand the forces he saw as “most likely to cultivate stable and common conceptions of reality:”

² Pooley goes on to detail these underlying motivations: “Shils had his own intellectual reasons for narrating the history in the manner that he did (first in 1948 and again, with more clarity, in 1951)--reasons rooted in his evolving and deeply engaged search for the underpinnings of modern social order... In a sense, however, his reasons did not matter once the narrative itself was released to the American sociological public; Lazarsfeld and Katz had their own reasons for adopting the historical picture that Shils put forward--reasons largely centered on scholarly competition and norms of originality. The powerful-to-limited-effects narrative in *Personal Influence*, in turn, was so widely embraced in the late 1950s for a still different set of reasons--because of the scholarly support it lent to the public intellectual defense of American popular culture, in the context of an evolving cold war liberalism... The staying power of this limited-effects narrative was ultimately guaranteed, however, by the newly institutionalized, would-be discipline of ‘communication’--which retained the story line as a usable, and teachable, past” (Pooley, 2006b, p. 135).

the “overall pattern of programming to which total communities are regularly exposed over long periods of time,” and in particular the effects of television, “the primary common source of socialization and everyday information (mostly in the form of entertainment) of otherwise heterogeneous populations” (Gerbner, 1998, pp. 179, 177). McLuhan, as I have already discussed in detail, saw the content of media as a mere distraction, and sought the true ‘message’ of mediation in the way it altered our patterns of thought and interaction. I return to McLuhan’s thoughts on Lazarsfeld below, but first I briefly examine the methods and findings of the Cultural Indicators project, one of the most central and long-running endeavours within communications theory to understand the effects of media on cognition.

As first proposed by Gerbner and Larry Gross, then subsequently developed in a series of ‘Violence Profiles’ in the *Journal of Communications* with several collaborators, this project accepted rather than rejected Lazarsfeld’s view that isolated propaganda or advertising campaigns were relatively ineffective in persuading people to “alter conventional conceptions, beliefs, and behaviors” (Gerbner & Gross, 1976, p. 181). Instead, they sought to situate the effects of media content in relation to the ‘construction of social reality’ (cf. Berger & Luckmann, 1990; Searle, 1995) through the broader notions of cultivation and enculturation. Media content may not cause us to change our minds about specific points in many cases—whether voting for a different political party or buying a different laundry detergent—but Gerbner and his co-authors argued that television in particular played a demonstrable role in the production of our ‘conventional conceptions, beliefs, and behaviours’ about how the world works. Working from the assumption that “the environment that sustains the most distinctive aspects of human existence is the environment of symbols,” and that this environment is being constructed in particular ways by an “increasingly professionalized, industrialized, centralized, and specialized” set of institutions, they studied television as the “chief

source of repetitive and ritualized symbol systems cultivating the common consciousness of the most far-flung and heterogeneous mass publics in history” (Gerbner & Gross, 1976, pp. 172–173). They operationalized this quasi-anthropological premise through a two-step process. First came content or “message system” analysis, to uncover whether the content of media systematically follows certain patterns—particularly in terms of its depiction of violent acts and of sexuality, the two most enduring foci of concern in media content—then survey research was undertaken to determine whether our perception of the world matched up more closely with the ‘symbolic environment’ portrayed on television, or with the frequency of events in the real world (Hawkins & Pingree, 1981). The Cultural Indicators group’s published findings uniformly indicated the former was the case, and they were eventually distilled into a “Mean World Index.” Since the frequency with which crime is depicted on television far exceeds the likelihood that one may in fact be a victim of crime, they hypothesized and found evidence that television viewing cultivated not only an increasingly strong fear of crime, but moreover a higher score on key indicators of ‘interpersonal mistrust.’ Television, it seemed, made us believe in a ‘mean world,’ more strongly as we watched it more often (Gerbner, Gross, Morgan, & Signorielli, 1980).³

As the results in this table suggest, the perception of a ‘Mean World’ is not something that is directly communicated by the content of any particular program, but an understanding of social reality cultivated over a long term, revealed through statistical analysis and mediated by a range of other social factors (including sex, ethnic minority status, educational attainment, and use of other media such as newspapers). Even when controlling for all of these, however, the ‘Cultivation

³ The key indicators, specifically, forming the ‘Mean World Index’ asked survey respondents to agree or disagree whether “most people are just looking out for themselves,” “you can’t be too careful in dealing with people,” and “most people would take advantage of you if they got a chance.” Television viewing “overall is significantly associated with the tendency to express mistrust ($r = .12, p < .001$)” (Gerbner, Gross, Morgan, & Signorielli, 1980, p. 17).

	Percent saying that you must be careful in dealing with people ^a			Percent saying that people are selfish rather than helpful ^b		
	New Jersey (N = 420) ^c			New Jersey (N = 413) ^e		
	% Light viewers ^c (N = 97)	CD ^d	gamma	% Light viewers ^c (N = 101)	CD ^d	gamma
Overall	52	+10	.21*	56	+ 8	.17*
controlling for:						
Sex						
Male	56	+10	.21	62	+ 2	.05
Female	48	+12	.22*	50	+13	.27*
Grade in school						
7th	49	+16	.32**	56	+ 3	.06
8th	54	+ 5	.10	56	+13	.28*
Ethnic group						
Ethnic	62	- 2	-.03	60	+ 7	.16
Non-ethnic	48	+14	.28**	53	+ 9	.19
Newspaper reading						
Every day	49	+16	.31*	59	+ 4	.10
Sometimes	55	+ 6	.12	54	+10	.20
Network news watching						
Almost daily	53	+13	.26	57	0	-.00
Once in a while	53	+ 8	.16	57	+ 5	.12
Hardly ever	48	+13	.26	52	+18	.36*
Father's education						
No college	56	+12	.23	60	+ 5	.10
Some college	48	+ 8	.17	54	+ 3	.06

* $p \leq .05$ (tau)

** $p \leq .01$ (tau)

^a "Can most people be trusted, or do you think that you can't be too careful in dealing with people?"

^b "Would you say that most of the time people try to be helpful, or that they are mostly just looking out for themselves?"

^c Percent of light viewers giving the television answer

^d CD (Cultivation Differential) is the percent of heavy viewers minus the percent of light viewers giving the "television answer."

^e Total number of respondents

Figure 26 - Cultivation analysis summary

Differential' remains positive, implying that heavier television viewers tend to be more fearful and less trusting of others. The 'Violence Profiles' dealt with violence in a rather different fashion than many other projects:

Conventional wisdom and fearful people, themselves victimized by images of violence around them, might stress the one or two in a thousand who imitate violence and threaten society. But it is just as important to look at the large majority of people who become more fearful, insecure, and dependent on authority, and who may grow up demanding protection and even welcoming repression in the name of security (Gerbner, Gross, Signorielli, Morgan, & Jackson-Beeck, 1979, p. 196).

The social stakes of violent media, for cultivation analysis, are not that some children or adolescents may imitate violent media content—as in the classic 'Bobo doll' studies of Albert Bandura (Bandura, Ross, & Ross, 1961), which have become paradigmatic in the analysis of violent video games (Anderson & Bushman, 2001). While this may be a minor risk, of far greater concern is that the cultivation of a set of beliefs corresponding to a 'mean world' may lead to troubling consequences for democracy. A fearful populace is one more likely to not only 'demand protection' but 'welcome repression' in pursuit of safety and security. Along similar lines, other analyses noted that heavy

television viewers (particularly of crime dramas) were more likely to report distinct patterns of action as well as belief, obtaining locks, guns, or dogs for personal protection in greater proportions (Gerbner, Gross, Jackson-Beeck, Jeffries-Fox, & Signorielli, 1978, p. 199). A ‘final note of gloom’ observed using general social survey data was that, as compared with light viewers, heavy viewers were more likely to envision another war within the next ten years (*ibid.*, p. 202).

The distinctive approach of cultivation analysis meant that McLuhan, who typically scorned ‘content analysis,’ regarded Gerbner and his group as potentially valuable allies. He certainly had no time for the dogma of ‘limited effects,’ for which he offered the backhanded *apologia* that “Professor Lazarsfeld’s helpless unawareness of the nature and effects of radio is not a personal defect, but a universally shared ineptitude” (McLuhan, 1994, p. 299). Here McLuhan specifically cites Lazarsfeld’s assessment that Hitler “did not achieve control through radio but almost despite it, because at the time of his rise to power radio was controlled by his enemies.” His own assessment, by contrast, was that the German populace of the 1930s “danced entranced to the tribal drum of radio that extended their central nervous system to create depth involvement for everybody” (*ibid.*). In other words, it was some direct neurophysiological effect of the radio, irrespective of whether the content transmitted was favourable to Hitler, which created sociopolitical circumstances conducive to the rise of National Socialism. For McLuhan, this was the ‘hot’ nature of the radio medium, transmitting a stream of information in relatively high definition through one sensory modality.⁴ Cultivation theory likewise suggests an understanding of media effects which differs from Lazarsfeld’s focus on specific informational content. In fact, it seems a necessary complement to McLuhan’s rather

⁴ Elihu Katz expands on this and another ‘virtual debate’ between Lazarsfeld and McLuhan, noting that in primary material like the diaries of Victor Klemperer from 1930s Germany, we may find evidence for both views: that of Hitler’s voice on the radio as “the Redeemer” come to “the poor” in the potent “language of the Gospels exactly,” and that of the minimally affected audience ‘barely listening’ in a truck stop restaurant (Katz, 1999).

speculative account of radio's effects on the nervous system: that the role played by the media in the rise of a movement like National Socialism could not be fully accounted for by the specific contents of propaganda broadcasts by, for, or against Hitler, but rather the media cultivated a broad understanding of Germany as under threat, betrayed, and in need of security. These symbolic framings predated and eventually furnished the raw materials for an increasingly powerful Nazi propaganda machine, which even Lazarsfeld admitted was likely to have affected public opinion through simple monopolization of distribution channels (supplemented, of course, by 'primary' associations, face-to-face contact in rallies, with brownshirts in one's community, *etc.*). The approaches of McLuhan and Gerbner are thus ultimately compatible, and present a pair of distinct alternatives to the 'canonical' limited-effects model of media. Both recognize that the media are not vectors for injecting packets of informational content into the minds of a passive audience, but that they nevertheless profoundly affect how we think and act.

They corresponded extensively over the course of their careers, most of all around the editing of essays that McLuhan had written for the *Journal of Communication* while under Gerbner's editorship. Gerbner was tapped in 1968 to write a biography of McLuhan for the *Encyclopedia Americana*, just as the latter was becoming a popular-culture phenomenon. On the recommendation of Merrill Panitt, then-editor of *TV Guide*, McLuhan also sought to foster collaborations between his Centre for Culture and Technology at the University of Toronto and Gerbner's Annenberg School of Communications at the University of Pennsylvania, researching for instance the differing "sensory profiles" of populations in Toronto long since extensively exposed to television, and a population in Athens before television was widely available there (McLuhan to Gerbner, December 29, 1965; a project which never saw fruition). Their correspondence indicates McLuhan's increasing interest in the neurobiology of communications media and new tools for studying this topic, along

with his efforts to persuade other important media researchers like Gerbner to share his enthusiasm. McLuhan was fascinated, for instance, with experiments run by Herbert Krugman of General Electric using head-cameras and electroencephalography to examine “whether it was the medium or the ‘content’ to which [subjects’ brain-waves] responded,” evidently seeing them as an extreme confirmation of his own thesis, and grounds for a flurry of speculation: “under electric conditions man has become angelic, disembodied information, instantly translated and transported. Naturally, I am working steadily on the electric angelism of man in the Magnetic City” (McLuhan to Gerbner, February 19, 1971). Gerbner, for his part, was “not so impressed with the brain waves of Krugman,” and contended that humanity’s becoming more ‘angelic, disembodied, etc.’ only makes us “all the more real, as we objectify ourselves in shared symbolic representations and mass communications” (Gerbner to McLuhan, March 9, 1971).

Regardless of Gerbner’s cool reception, McLuhan’s enthusiasm for new scientific and technological tools for measuring our interactions with media only increased over the years. By 1976 he was arguing even more clearly in his correspondence than in his published writings that all of the conceptual pairs for which he was best known—the written and the oral, the visual and the acoustic, the hot and the cool, figure and ground, *etc.*—were best understood as mapped onto the two hemispheres of the brain. The literate and industrial world was long structured according to the ‘left hemisphere’ traits fostered by written, visual media, while the “Third world—the world without the phonetic alphabet” had always been associated with aural/oral culture and the ‘right brain.’ All of this was disrupted, however, by the “new electronic milieu or environment which automatically pushes the right hemisphere into a more dominant position than it has held in the Western world since the invention of the phonetic alphabet” (McLuhan to Gerbner, September 3, 1976). Those most highly educated within the traditional Western university system—scientists, engineers, and

others most directly responsible for the development of electronic media—were in McLuhan’s view “left hemisphere people,” whom it so happens were “completely out of touch with the results and the formal characteristics of their own new electric technologies” (*ibid.*). But at the same time, scientific insights into brain function seemed to him the most persuasive and apt way to characterize the effects of media, and he pursued over many years a project with W. Arthur Hurst, an optometrist, to measure the visual habits of the “TV generation” using head-mounted cameras (McLuhan to Gerbner, August 30, 1977).

Despite (or perhaps because of) his intense fascination with these ideas, McLuhan’s comments regarding the brain and neurophysiology are some of his most dubious, and most difficult to write off as mere ‘probes.’ While in *Understanding Media* he charted a very different course from the traditional ‘harm to children’ framing of media effects, he adopted just this position in a letter co-authored with Hurst to the Governor of California, arguing for both “curtailment of TV viewing for children,” and “in its place, a planned ‘motor’ program, involving multi-sensory motor experiences for the child.”⁵ In his last letter to Gerbner he contended that “these discoveries amount to saying that TV is a kind of disease, physiologically considered” (McLuhan to Gerbner, September 19, 1977) for its effects on the patterns of visual perception; the latter was less than enthusiastic, though positive as always, noting that this material “seemed a little far out and physiological for our readers,” but that he would be glad to read a finished manuscript (Gerbner to McLuhan, October 4, 1977).⁶

⁵ McLuhan sent a copy of this letter, dated August 4, 1977, to Gerbner the following day (McLuhan to Gerbner, August 5, 1977). All correspondence between these authors was retrieved electronically from the George Gerbner archive at the University of Pennsylvania.

⁶ McLuhan and Hurst’s argument (based on head-camera data) was specifically that television reduced the perceptual ability to focus and coordinate motor activity around ‘near points’ within the gaze: TV, an “all-involving, isolated ‘sensory only’ instrument,” a cool medium in McLuhan’s terminology, “dissociates the sensory from motor action, creating ‘half a man’” (McLuhan and Hurst to Gov. E.G. Brown, August 4, 1977). The gendered phrasing here was not incidental, as they gave this issue a distinct spin that is echoed in contemporary debates around ADHD. Learning

No finished manuscript was ever produced based on this work with Hurst, though the latter continued to promote the importance of his research with McLuhan, and his program of sensorimotor exercises to counteract the negative consequences of television. At one point Hurst testified to the Ontario Legislative Assembly that television was a key contributor to the rise in learning disabilities, and that federal funding was needed to counteract its effects (Cordiano & Hurst, 1994, sec. Standing Committee on Public Accounts). These efforts notwithstanding, the neurobiology of communications media has long remained a minor field of research on many levels. It is, first of all, rarely remembered as a prominent element in McLuhan's later work. This is likely in part because it never received a full published treatment before his death, and in part because his claims for strong hemispheric specialization are seeming increasingly problematic, both in themselves and in the binary cultural and gendered implications he drew from them (McLuhan, 1978).⁷ Moreover, the approach to media studies pioneered by McLuhan—concerned with formal properties of technologies rather than the content of mass media—has never really taken off. There have been periodic attempts to revive it (Levinson, 2001), and his quips are still routinely cited by popular writers, but this approach has long been the “least common form of media analysis” (Meyrowitz, 1998, p. 106). On the whole, media and communications studies have not really embraced McLuhan's call to examine the neurobiological effects of new technologies, nor have these effects become a major focus within the cognitive sciences. In the next section, I consider what research has since been carried out along these lines, and raise some points of critique. I wish to emphasize, however, that these are intended in the spirit proposed by the ‘critical neuroscience’

disabilities, specifically dyslexia, were said to be “almost entirely a male characteristic,” driven by heavy TV viewing and failure to develop the fine ‘near-point’ motor skills which “automatically develop in female activities (skipping, cooking, grooming, sewing, etc.)” (*ibid.*).

⁷ With regard to the idea of hemispheric specialization in itself, fMRI data strongly suggests there is no “whole-brain phenotype of greater ‘left-brained’ or ‘right-brained’” individuals (Nielsen, Zielinski, Ferguson, Lainhart, & Anderson, 2013).

movement, as calls to deepen and extend rather than undermine scientific research (Choudhury, Nagel, and Slaby, 2009). There are problems with McLuhan's own framing, and with many contemporary claims that new media are 'changing our brains,' but these only confirm that we need to better understand the interactions between technology, neurobiology, and consciousness.

Toward neuromediation.

Where do we stand now with respect to the neurobiology of communication? Have we advanced considerably beyond the antiquated notions of hemispheric specialization and 'sense ratios' as developed by McLuhan? In brief: yes, we have advanced, but no, not considerably. In this section I begin by considering what research has been done in the cognitive sciences examining media effects, the findings of which may be separated into two straightforward categories: those implying positive, and those implying deleterious effects on the mind. In parallel with my discussion of this research, I discuss how in turn these findings have been presented *within* the popular press, and criticisms that have been raised against them. After evaluating some of the most representative and significant findings on either side, I examine some alternatives to this dichotomous framing. In the final section, I then consider two ways in which cognitive science research goes beyond an explicitly stated interest in understanding media effects: first in its pervasive use of media technologies and then, perhaps more ominously, in new initiatives to not only represent but simultaneously *intervene in* the functioning of the plastic brain.

I begin with the research which seems most evidently to support the concerns of those who panic that new technologies are 'changing our brains.' What data supports the contention that "young people in particular are becoming unempathic, passive, intellectually shallow and uncritical, desensitized, depressed, and attention deficient because of cyber technology" (Choudhury &

McKinney, 2013, p. 194)? Nicholas Carr has argued that scientific evidence for this idea is substantial and mounting (Carr, 2011); Choudhury and McKinney suggest quite the opposite. Cognitive neuroscientist Kathryn Mills contends in a review that while the World Wide Web has changed “our way of interacting with each other and our collective history,” requiring new skills which are likely to be reflected in our brains somehow, “there is currently no evidence to suggest that Internet use has or has not had a profound effect on brain development” (Mills, 2014, p. 387).

Hence there is little in the published literature to satisfy the likes of Carr and Susan Greenfield, or those who wish to put their alarmism definitively to rest. Like all of these commentators, I am concerned here primarily with the effects of new digital technologies. A more thorough case for their distinctiveness can be elaborated, however, on the basis of the foregoing analysis. First, as outlined in the previous chapters, the rise of digital tools is intimately connected with the rise of cognitive science itself, and so it is of particular interest for this project how the discipline has looped back to investigate the cognitive effects of these very technologies. However, the importance of digital media in current public debates stems not only from its novelty—moral panics around media effects always tending to focus on the newest forms—but from the same universality which made it so appealing to cognitivism. Although they have introduced new possibilities, computing and the World Wide Web are not radically different and distinct from old media. Rather, the change that comes with digital technology stems from its capacity to gather together and facilitate rapid access to all prior forms of symbolic mediation from across history. Hence the same networked medium through which I have conducted the research and composition of this document is also my means of accessing television content, recorded music, interactive games, and any number of other forms of representation, all now translatable into digital code. I can retrieve at will the cave paintings at Lascaux or examples of Sumerian cuneiform, in a multiplicity of

static and dynamic visual representations. All this is to say that our mode of interaction with these older media has been substantially altered by the rise of this new medium—and that along with everyone else labouring on computers, I am perpetually confronting possibilities for distraction.

Where once the primary worry seemed to be television and its cultivation of passive viewership, concerns have now shifted to overstimulation and hyperactive patterns of attention (Hayles, 2007; Stiegler, 2010). Only a few decades ago, McLuhan and Hurst argued that learning problems were caused by television's decoupling of motor activity from vision at the 'near point,' now we are as likely to hear arguments that they stem from excessive video gaming, which does precisely the opposite by rendering the screen interactive. Carr's McLuhan-influenced writings weave together historical reflections on technology with neuroscientific research to contend that the hyperlink-based and highly interactive nature of the Web is turning users into 'cursory readers,' 'hurried and distracted thinkers,' and shallow people. This is in effect a milder version of the argument which I consider in the next chapter: that short of literally causing attention deficit disorder, our interactions with computers are making us all distracted and unable to focus on single tasks for very long (like reading or writing lengthy books, for instance). It is certainly an argument that resonates well with anyone who has found themselves distractedly, almost unintentionally, browsing some unrelated pages for far too long in the midst of working on a project.

It is buttressed, moreover, by the now-ubiquitous metaphor schema of mind/brain-as-electronic-system: though elsewhere he recognizes the irony of this trope for his argument, Carr contends that "if, knowing what we know today about the brain's plasticity, you were to set out to invent a medium that would rewire our mental circuits as quickly and thoroughly as possible, you would probably end up designing something that looks and works a lot like the Internet" (Carr, 2011). The Internet, he argues, offers just the kind of cognitive stimuli – 'repetitive, intensive,

interactive, addictive' (*ibid.*) – which seem most likely to alter our neural functioning. The argument holds as well for electronic games and other applications of computer technology. Despite the spreading perception that we may be acquiring a technologically-induced “continuous partial attention deficit disorder,” others have argued this is either a fiction or a temporary adaptation, invoking neuroscience equally to argue that these same ‘hyper’-interactive media are enhancing our fluid intelligence and sensorimotor capacities (Casio, 2009; S. Johnson, 2005). Unfortunately much of this commentary, including Carr’s, engages with science in a way that seems to bear out his view, covering a wide range of basic research in an expansive, speculative, compelling, but ultimately fairly superficial fashion. What additional light can we shed on these debates by attending more closely to the processes by which this research is produced? Is Google making us stupider, shallower, smarter, or something else entirely? How do we know?

The evidence produced by the cognitive sciences regarding media effects is of a few different types. There are traditional survey and observational measurements of populations, seeking to analyse correlations between media use and particular psychological traits; there are studies which directly measure some physiological property, be it through brain imaging, electroencephalography, gaze-tracking cameras, or some combination of techniques; finally, there are a handful of studies employing animal or computational models. The first of these are certainly the most numerous, dating back to concern with culture and language-acquisition in early cognitive psychology. As one of the major contributors to the cognitive science of media, Patricia Greenfield (not to be confused with Susan), argued in a paper with Jerome Bruner, an underlying premise of this research is that “intelligence is to a great extent the internalization of the ‘tools’ provided by a given culture.” Thus we could improve our understanding of both domains—intelligence and culture—by “comparing intellectual development in cultures with radically different technologies” (P. M. Greenfield &

Bruner, 1966, p. 90), and by examining the psychological consequences when new forms of media are introduced to previously unexposed populations, as Greenfield reviewed in her own McLuhan-influenced volume (P. M. Greenfield, 1984). Among the most significant results she flagged were those relating to violent content, and this remains an active focus of psychological research and public concern today.

Researchers proposed that even early video games such as *Space Invaders* caused increases in aggression among children (P. M. Greenfield, 1984, pp. 103–104). More recent studies have continued to build on this consensus, arguing that despite other consequences increased visual realism in games has done nothing to alter the basic conclusion that violent content promotes aggressive cognition. What does this really mean, though? In the context of one particular study, a selection of Nintendo video game consoles ranging from the Super Nintendo (1990 – low resolution and realism) to the Wii (2006 – higher resolution and visual realism) were employed, alongside games from the *Mortal Kombat* series to represent the violent condition, and baseball games, to represent a similarly ‘exciting’ but non-violent stimulus. Subjects were asked to play the games for 15 minutes, and then complete a heart rate measurement, a State Hostility Scale, and a Word Completion Task. The SHS asks participants to rate statements on a scale of agreement, representing hostile cognitions: *ie.*, “I feel like banging my head on a table,” (1 – *not at all*) or (5 – *extremely*)? The Word Completion Task gives word fragments such as “K I _ _,” which could be completed in either an aggressive fashion (KILL, KICK), or nonaggressively (KISS, KILT, *etc.*). Finally, they were asked to rate their own sense of immersion in the games (Barlett, Rodeheffer, Baldassaro, Hinkin, & Harris, 2008, p. 554). From a series of statistical analyses, the researchers concluded that immersion was positively correlated with visual realism, but that neither immersion nor realism “seem to matter much” (*ibid.*, p.560) in terms of the overall association between violent

games and hostility, and so the effects of violent game content on aggression are independent of how detailed the world is or how immersed the player feels within it. Other studies have been carried out according to fundamentally similar methods, but finding opposite results, that while increasingly realistic video games did correlate with higher levels of reported immersion and recorded physiological arousal, there was no significant link between violent games and aggressive cognition in either case (Ivory & Kalyanaraman, 2007).

While not universally supported, there is a building consensus for limited but significant effects of violent media content in promoting violent or hostile cognitions. The most important method for establishing such consensus within the research community is the meta-review (eg. Anderson & Bushman, 2001), which aggregates many prior studies, typically weighting experimental designs more heavily. Video games have remained a primary focus in this regard, with several analyses arguing there is clear support for their cognitive effects (Anderson et al., 2010; Anderson & Bushman, 2001). Others have produced their own meta-analyses contending, however, that whatever measures we may devise to show increases in subjective hostility, the influence of violent games “on serious acts of aggression or violence is minimal, and publication bias is a problem in this research field,” with researchers tending to use “problematic unstandardized measures of aggression” to produce their strongest results, while the best measures of aggression produced the weakest effects (Ferguson & Kilburn, 2010). The former viewpoint seems to be more widely espoused within the popular press, to be sure, but also within the professional literature—the cross-cultural 2010 meta-review by Anderson *et al.* has been cited roughly five times as often as Ferguson and Kilburn’s null result.⁸

⁸ The point they raise about publication bias is a valuable one, particularly when considering meta-reviews, and in the concluding section of this project I discuss how an online grassroots movement of sorts has arisen to take the scientific publication process to task, particularly in the field of psychology. It is worth noting as well that the social stakes of these

One of the other interesting ways in which psychological research has begun to consider the potential relationship between new media and hostile cognition is in relation to Internet ‘trolling.’ This form of conduct, dating back to the days of Usenet, is characterized by a range of deliberately offensive, confrontational, and threatening activity directed against others online for the perceived humour value of outraged responses. Some commentators, alluding to Plato, have called this a ‘Gyges effect’ (Hardaker, 2013), emerging from the pseudo-anonymity of online interactions; psychologists, for their part, have sought to uncover correlations between trolling behaviour and established personality traits, particularly the ‘Dark Tetrad’ of sadism, psychopathy, narcissism and Machiavellianism, finding that “trolls are prototypical everyday sadists” (Buckels, Trapnell, & Paulhus, 2014). It seems an increasingly important question whether the formal properties of the Internet play some role in the genesis of such conduct, or rather just provide greater visibility to a personality type that has always existed at roughly the same rate.

None of this, however, deals directly with the brain. As the authors of one of the few relevant imaging studies recognized, “Little is known about the relation between media violence and brain functioning” (Mathews et al., 2005, p. 291). The ubiquity of neuro-discourse within these debates belies the fact that studies of the traditional psychological type are by far the most numerous investigations of media effects, greatly outnumbering those which investigate the brain by measurement or modeling. To expand on the observation that evidence in these debates is ‘thin and ambiguous’ (Choudhury & McKinney, 2013), this psychological evidence is substantially thicker than the properly neuroscientific data, but equally ambiguous. Thus it serves a crucial function in

issues are right on the surface within the papers: Anderson and Bushman open their first meta-analysis by invoking a series of recent school shootings, and the shooters’ affinities for video games, while Ferguson & Kilburn conclude by arguing that “Psychology, too often, has lost its ability to put the weak (if any) effects found for VVGs on aggression into a proper perspective. In doing so, it does more to misinform than inform public debates on this issue” (Ferguson & Kilburn, 2010, p. 177).

supporting arguments which draw broader conclusions from fairly sparse data, while nevertheless couched in the language of neuroscience throughout. Nicholas Carr's book is far from the only recent volume to argue that digital technology is 'rewiring our brains,' and other treatments rely more heavily on the media-violence literature (Small & Vorgan, 2009).

Turning now to those studies which *do* specifically address the brain and media, while overall they remain few and far between, the effects of violent games are one of the most commonly considered areas in this regard (Mathews et al., 2005; Mathiak & Weber, 2006; Rene Weber, Ritterfeld, & Mathiak, 2006). These recent papers are directly positioned as novel extensions of prior psychological research finding a link between violent games and aggression, particularly in children, and their inspiration is drawn from a 'communibiological paradigm' (Beatty, McCroskey, & Heisel, 1998). This proposed paradigm for communications research understands all communication, both face-to-face and technologically mediated, as rooted in neurobiology, particularly in the "circuitry involved in attentional focus and regulation,"




Figure 2.

Explorative mapping of active phases as compared to the passive condition (virtual death). **A:** The save condition (left column) elicited higher activity (orange overlay) at motor (cerebellum, supplemental motor area (SMA), and left-lateralized motor cortex) and executive areas (dorsolateral prefrontal cortex) in addition to lower activity (blue overlay) in bilateral temporoparietal junction. **B:** The active fighting condition (middle column) was linked to distributed inhibition of frontal, limbic, auditory, and right motor areas and activation of cerebellar, visual, temporoparietal, and cingular areas. The under attack (**C**) and the potential danger (**D**) conditions activated only in motor related areas (under attack: cerebellum, SMA, and motor cortex; potential danger: only SMA).




Figure 3.

Single-subject analysis with time course of exemplary raw data at the rACC (**A**) and the amygdala (**B**) as compared to violent scenes (marked with red lines). Regardless of the noise, dips after each violent event are clearly delineated. **C:** Overlay of BOLD signal changes at the level of the amygdala of a single subject on its EPI data. The (unsmoothed) EPI shows a low amount of dephasing and the overlaid activation is clearly located within the right amygdala in this subject.

Figure 27 - Brain images of media violence exposure

and the 'behavioural inhibition system' (*ibid.*, p. 208-209). While the original formulation of this paradigm suggested that the "origin of the brain structures and circuitry" involved were "inborn characteristics of individuals" (*ibid.*, p. 211), more recent neuroscience in this vein seeks to consider the effects of communication on the plastic brain. A major region of interest in the case of violent games is the amygdala, a brain area previously implicated in emotion, particularly aggression and fear (Rene Weber et al., 2006; Coccaro, McCloskey, Fitzgerald, & Phan, 2007). These studies employ brain scanning paired with coding of the different phases involved in violent game play (Mathiak & Weber, 2006, p. 952). Groups of coders, typically graduate students, view the subjects' actions within the game world and code different segments as containing violent events, active fighting, player death, and so on. fMRI or PET information about the flow of blood within the brain is gathered concurrently, mapped onto a standardized brain 'atlas' (either the Talairach space or the Montreal Neurological Institute system), and then the researchers look for elevated activation, as measured by the average oxygenation or glucose consumption within particular sets of 'voxels' or regions of interest in the subjects' brains. Data are then 'smoothed' and the relative activations of different brain areas are given graphical representation on 'heat map' diagrams like those in Figure 27 (Mathiak & Weber, 2006, pp.952, 953). These studies all tend to find confirmations for their hypotheses: that violent events will be correlated with peaks in amygdala activation, and perhaps just as importantly, that individuals respond in violent games much as they would to other sorts of stimuli, hence having a subject play video games within a scanner "can be considered an ideal paradigm to study behavior in immobilized subjects" (*ibid.*, p. 948).

As several prior studies have argued, images produced by fMRI and PET brain scanning are far from neutral, direct, or unmediated evidence (Beaulieu, 2001, 2002; Dumit, 2003). They are images constructed in complex ways, impossible to produce without the mediation of computerized

statistical analysis. Much like the DENDRAL system discussed previously—and the visions of Babbage and Licklider which inspired it—this is an interactive process in which human perceptual capacities are simultaneously extended by and used to direct a semi-autonomous technological system, fitting the data to standardized maps of anatomical areas, while selecting different regions of interest, mathematical operations, and software packages to suit their project.⁹ This means, of course, that there is considerable room for ‘fudging’ results to fit hypotheses, and ensuing questions about how well claims of localization in specific brain areas across subjects can really be substantiated. It has been argued that the weaknesses of brain imaging techniques in this regard mean we should discard them entirely, at least for the time being, in favour of more proven psychological and psychophysiological methods (Uttal, 2003, 2011).

The community of neuroimaging researchers is quite attentive to these issues, however, and nearly every paper is laden with caveats and reflexivity about its own limitations.¹⁰ While they contended, for instance, that the amygdala signal changes were “surprisingly consistent,” with high effect sizes (Mathiak & Weber, 2006, p. 954), all studies noted that they could not prove whether the

⁹ Nicholas Carr’s more recent work has turned to a critique of this very process of automation as extended to human cognitive faculties, arguing against contentions that increasing automation would ‘free’ the human mind from mundane tasks and allow it to achieve ever greater things. He contends that “Automation is different now... Many software programs take on intellectual work—observing and sensing, analyzing and judging, even making decisions—that until recently was considered the preserve of humans. That may leave the person operating the computer to play the role of a high-tech clerk—entering data, monitoring outputs, and watching for failures. Rather than opening new frontiers of thought and action, software ends up narrowing our focus. We trade subtle, specialized talents for more routine, less distinctive ones” (Carr, 2013).

¹⁰ One noted for example that although many would consider the media coding aspect of this research to be clearly the less reliable portion, intercoder reliability for that portion was fairly high, and in fact “the necessary data normalization and transformation (such as standardizing the individual participant’s brain to the Montreal Standard Brain), spatial filtering and smoothing, hemodynamic response folding, and the selection of brain areas assumed to be representative for the ROI [region of interest] under consideration, make fMRI a less reliable measure than often assumed. Thus, the resulting brain patterns serve as the best possible measures for stimulus response, but they are not error free” (Rene Weber, Ritterfeld, & Mathiak, 2006, p. 52). It is worth mentioning as well that new techniques for ‘optogenetics’ show promise for more closely approaching the measurement of single neurons, using targeted light (Packer, Roska, & Häusser, 2013). Multi-voxel pattern analysis in functional imaging also implies a shift away from localizing function in specific anatomical regions to more broadly distributed networks, with both developments suggesting that critiques raised from within and outside the imaging community may be at least partially addressed by new technologies.

changes they observed would lead to increases in actual aggressive behaviour. They all further speculated that additional repetition of such stimuli *could* lead to enduring changes in brain structure, but recognized that their experimental designs with controlled and limited exposure could not shed light on this possibility. One study did attempt to evaluate these changes in an interesting, oblique fashion, however. While other papers considered only individuals without psychiatric diagnoses, this one specifically sought to compare fMRI scan data from a group of 71 adolescents diagnosed with a disruptive behavior disorder (DBD), with at least one “significant symptom of aggressive behavior toward people, animals, or property within the past 6 months,” and a control group without DSM diagnoses (Mathews et al., 2005, p. 288). They calculated for each subject a ‘Media Violence Exposure Index’ based on survey data about the adolescents’ television and game-related habits, then had them complete an unrelated ‘Counting Stroop [CS] task’ involving number patterns while inside a fMRI scanner. From the data obtained, “a Student *t* value was calculated for each pixel to create individual activation maps that were then transformed into Talairach space” (*ibid.*), leading them to conclude that, as hypothesized, subjects with high media violence exposure showed reduced frontal lobe activation “similar to that seen with aggressive individuals” (*ibid.*, p.291).

Mathews *et al.* did not consider the amygdala, as did the other study—noting that its specific scanning technique and CS task did not allow for good measurements in that particular brain region—instead building upon prior work showing changes in frontal lobe regions associated with aggressive behaviour. The particular regions of interest in this study were the medial and inferior frontal gyrus, and the anterior cingulate cortex, all of which are implicated as well in the broader functions of attention and executive control. The entire group with diagnosed DBD, nearly half of whom were diagnosed with comorbid attention deficit disorders, displayed activation only in the medial frontal gyrus, and neither of the other regions; the control group with low media violence

exposure according to the survey instrument displayed activation in all three areas; the control group with high media violence exposure, by contrast, showed similar patterns as the DBD group, with activation only in the medial frontal gyrus. “Taken together,” they conclude, “these findings suggest that media violence exposure may have an influence on brain functioning whether or not trait aggression is present” (*ibid.*). Reduced anterior cingulate cortex function is in turn commonly associated with ADHD, with the implication it may lead to reduced capacity for both attentional and emotional regulation. This study hence links media violence not only with aggression but with a broader complex of deficiencies in executive function, albeit in a highly preliminary fashion. These may give rise to a range of behavioural problems, and so this seems a fruitful direction for further research.

Most of these studies feature the heat-map images which have become paradigmatic tokens of persuasion in contemporary neuro-discourse, another compelling bit of evidence to be recruited in service of the media-aggression link. But it is worth recalling that such images do not straightforwardly document the truth of the brain—they are data ‘made presentable,’ as Dumit contends (Dumit, 2003)—and at the same time that their constructed nature is by no means hidden or downplayed within the professional literature. On the contrary, documenting the process of construction and presentation is the major concern of methods sections within these publications, and some of the sharpest critiques of dubious techniques for proving correlations by brain imaging have come from within the research community.¹¹ The issue, as cognitive scientists, journalists, and

¹¹ The most entertaining episode in this regard was entitled “Neural correlates of interspecies perspective taking in the post-mortem Atlantic Salmon: An argument for multiple comparisons correction,” a paper eventually published in the *Journal of Serendipitous and Unexpected Results*, which used standard statistical manipulations to find ‘significant BOLD signal change’ in a cluster of voxels within the brain cavity of a dead salmon purchased from the grocery store (Bennett, Baird, Miller, & Wolford, 2011). The conclusion they drew from this obviously false result was that better corrections were needed and not, it is worth emphasizing, that functional neuroimaging research was a hopeless endeavour. Along similar lines, a manuscript which once circulated informally as “Voodoo Correlations in Social Neuroscience” has recently been published arguing for ‘Puzzlingly High’ effect sizes in a sample of 55 imaging articles: its thesis is that the *r*-values for

many others have recognized, is that these nuances and caveats often do not travel intact with the results as they diffuse into the public sphere. In the mass media “elucidating the neurobiological correlates of a phenomenon” is regularly “presented as comprising a full explanation for its existence,” though the “actual explanatory power of the biological information alone” is more often “imperfect” and partial (O’Connor, Rees, & Joffe, 2012, p. 224). Neuroscience is almost inevitably recruited for persuasive effect in the debates I consider, but more conventional social-scientific research—alongside varying measures of opinion and conjecture—is just as essential to the production of a comprehensive argument from the sparse and underdetermined data.

In another sense, the question of media violence is itself on its way to becoming old news. Nicholas Carr is just one of many commentators who have applied the ‘changing our brains’ argument less to specific media content, however worrisome, and more to the formal properties of digital technology. Above all, concern is mounting with respect to the transformation of so many distinct spheres of our lives by networked computing. The fear is that these forms of mediated experience will grow to outcompete and replace more authentic modes of human interaction, with potentially dire consequences. ‘Addiction,’ whether to the Internet or to video games, is one prominent though controversial means of interpreting this process through the lens of pathology (Byun et al., 2009; Pies, 2009). This has become an increasing focus of public concern in the past decade, with a series of troubling incidents in which individuals committed suicide or murder over electronic games, parents neglected children to the point of their deaths, and one player even died of

correlations in excess of .8, as often found in these papers, are mathematically impossible given the limited reliability of both imaging and behavioural measurements. Just as running several studies and only picking out the confirmatory results can cause false positives, the typical process for computing activation within fMRI voxels can produce a ‘nonindependence error,’ thus inflating the reliability of observed correlations and even leading to possibility of creating a ‘signal’ from ‘pure noise’ (Vul, Harris, Winkielman, & Pashler, 2009, pp. 279, 284). This strongly worded ‘fake signal from pure noise’ argument closely resembles the rhetorical strikes of *Quest* in Collins’ classic analysis of the gravitational waves controversy (Collins, 1981), suggesting that concern with the agency of technology in scientific debates may not be far removed from the traditional sociology of scientific knowledge.

heart failure, apparently spurred on by exhaustion and dehydration after a marathon multi-day gaming session in South Korea (Kosoff, 2014).

This moral panic recruits a newly ubiquitous neuro-trope: the ‘dopamine squirt.’ A cursory examination of Google News reveals dozens of articles within just the past few months of writing, contending that the real secret to the latest flavour-of-the-month game—be it *FarmVille*, *2048*, or whatever may have superseded them at the time of your reading—lies in its harnessing of this particular neurotransmitter (Krisch, 2014). Much the same is typically said of interaction with social networking sites like Facebook, which equally play host to online browser-based games. Publications like *AdWeek* have described what elsewhere I have referred to as an ‘attention economy’ as instead a “dopamine economy,” a corollary of neuromarketing-speak which seeks to “engage consumers...at a brain chemistry level,” and which implies that the gaming industry may “hold the keys to the future of advertising” (Hicks, 2014). The connection to addiction, in turn, comes from the fact that dopamine release, particularly in the nucleus accumbens, is strongly associated with addiction to drugs of abuse like cocaine and amphetamine (Sulzer et al., 1995; Kahlig et al., 2005). Hence linking behavioural addictions to digital media with these more demonstrable physiological addictions is a crucial element in achieving closure around the view that these are ‘real’ addictions.¹² In the case of gaming, there are indeed a few studies which support this connection.

Among the first of these was a PET scanning paper which used a radiotracer to detect changes in the levels of dopamine released within the brain during a game. Subjects first were asked to look at an empty screen, then to drive a tank in a game environment, while the scanner measured

¹² Many parallels for this debate are present in the gambling addiction literature, which is far broader and better established, and which now often is concerned with forms like ‘video poker,’ materially indistinguishable from other types of electronic games and equally well suited to study within the confines of a brain scanner (Linnet, Peterson, Doudet, Gjedde, & Møller, 2010; Schüll, 2012; Anselme & Robinson, 2013). I return to the issue of ‘neuromarketing’ at the close of this chapter.

the radiotracer levels in various parts of their brains. This particular radiotracer is sensitive to dopamine, in that dopamine successfully ‘competes’ with it at the synapses, and so decreases in the binding of the former indicate converse increases in levels of extracellular dopamine. The gaming task produced results indicating such increases in the striatum—a larger brain region which encompasses the nucleus accumbens—at levels comparable to those following injection of amphetamine or methylphenidate (Koepp et al., 1998). This article was a short “Letter” in *Nature*, and was more concerned with the very technique for demonstrating “behavioural conditions under which dopamine is released in humans,” and illustrating their specific PET method (*ibid.*, 266), than with rigorously testing the effects of games. Apart from small numbers (only eight participants), the study lacked any real control other than the blank screen, hence it was impossible to dissociate the supposed effects of games from those of mere visual perception in complex, changing scenes. Yet it remains widely cited in the press and in popular books as evidence for the similarity of video games and addictive drugs on a neurobiological level (eg. Small & Vorgan, 2009; S. Greenfield, 2014; Messaris & Humphreys, 2006). A further difficulty with its conclusions is that the analogy with psychostimulant drugs can run in a different direction entirely, an issue of framing which comes to the fore in the next chapters. Rather than being seen as comparable to drugs of abuse, the role of amphetamines and related compounds in *treatment* for conditions like ADHD may be emphasized, thus instead it may be argued that suitably designed games provide an ideal learning environment—either for formally diagnosed children or for the hypothesized multitudes of quasi-ADD sufferers (Gee, 2005; S. Johnson, 2005). Hence the same imaging evidence may be recruited equally by

opponents of ‘addictive’ gaming and by advocates of broader ‘gamification’ in education and society.¹³

This underdetermination of wider social conclusions should not be taken, however, to imply that we confront pure uncertainty in considering the effects of games on the brain. Other more recent studies have sought to triangulate the above findings using fMRI to measure variations between regular game players and those unexposed to gaming. Likewise, they found higher volumes and differing patterns of activity within the striatum of those who played often, lending further support to the idea that these interactive media may induce brain changes broadly associated with dopamine and attention-governing systems in the brain (Kühn et al., 2011). Equally, however, it could be a mere correlation, that games are inherently more appealing to individuals born with thusly enlarged brain regions. What would be needed to furnish more decisive evidence in these debates would be long-term, prospective studies of the effects of digital media on the brain, which at present do not exist.¹⁴ Instead we can only infer changes from anatomical correlations with media use, or from functional changes following short-term exposures. Where the effects of interest are hypothesized to be negative, ethics guidelines may imply that these are the only kinds of imaging studies which can be carried out (a problem that does not apply in the case of positive effects such as ‘brain training’). The other means of going beyond this evidence is by the strategy which I have identified above as most closely associated with cognitive science proper, as opposed to cognitive neuroscience: modeling.

¹³ This is the notion of putting ‘game design elements’ to use in a range of contexts where they have not traditionally applied – creating software that applies ideas from gaming such as ‘levelling up’ or virtual badges for ‘achievements’ to the pursuit of goals in education, employment, self-improvement and elsewhere (Deterding, Sicart, Nacke, O’Hara, & Dixon, 2011).

¹⁴ Such studies do exist, by contrast, in the media violence literature, and by consensus constitute some of the best evidence in support of the link between violent content and trait aggression (O’Connor, Rees, & Joffe, 2012). Again, though, any linkages between these arguments and the brain are highly speculative.

One of the most prominent researchers into media effects over the past decades, Dimitri Christakis, has recently begun to investigate attention and media ‘overstimulation’ through the somewhat surprising method of animal modeling.¹⁵ A recent analysis with two collaborators developed the

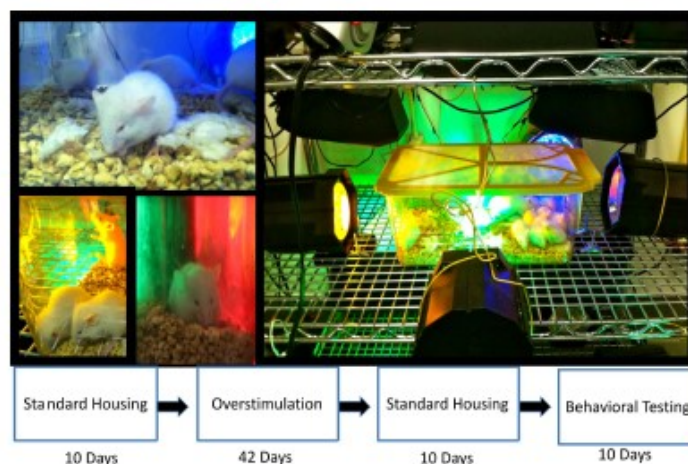


Figure 28 – Dimitri Christakis’ ‘mouse overstimulation chamber’

setup pictured here (Figure 28), randomly dividing 10-day-old mice into two groups, with the first being reared according to standard protocols as a control, and the second being treated identically but for 6 hours of nightly exposure to bright colours and audio from cartoons in the ‘overstimulation chamber’ (Christakis, Ramirez, & Ramirez, 2012, pp. 1–2). The mice were then subjected to four behavioural tests, examining how long they spent investigating novel objects in an enclosure as opposed to objects they had previously seen, their patterns of movement in an open field, as well as in a ‘Barnes maze,’ and an ‘elevated plus’ maze. ‘Overstimulated’ mice displayed patterns of behaviour that differed significantly from the controls, spending more time on the ‘open’ arms of the elevated plus, more time in the middle of an open field, and taking longer to solve the Barnes maze. From this the researchers inferred that the ‘overstimulated’ mice suffered important “deficits in cognition and behavior,” with “increased risk taking/anxiety, poorer short-term memory, and impaired learning” (*ibid.*, p. 3). How effectively this research demonstrates these effects depends, as with all animal research in psychology, on how well one believes these stimuli replicate the real

¹⁵ I discuss Christakis’ best-known paper, a longitudinal survey-based study suggesting a link between television exposure and attentional problems, at greater length in the following chapters, alongside media coverage of his research (Christakis, Zimmerman, DiGiuseppe, & McCarty, 2004).

phenomena of interest—namely television viewing for children—how well the behavioural tests operationalize the traits of anxiety, short-term memory, and learning, as well as one’s general opinion of drawing analogies from the brains and behaviours of mice to those of humans.

For Christakis, this fits well into a broader picture of television and digital media causing significant cognitive changes, above all in relation to attention, learning and memory. Such results do not stand on their own, but may be recruited alongside the other streams of research cited above in an attempt to triangulate the detrimental effects of technology. All this is again configured around a discourse of the ‘neurological adolescent,’ and potential for harm in critical periods of development. The overall pattern of Christakis’ research seems to be oriented toward demonstrating scientifically the truth of the proposition that all of us, but particularly our children, are being driven toward attentional problems by the media—be they diagnosed ADD, or similar but ‘subclinical’ symptoms. How strongly this thesis has been proven remains open for debate, but many have certainly interpreted it thusly. Christakis’ work is often seen as supporting the American Academy of Pediatrics guideline that children should not be exposed at all to television or other ‘screen-based’ media before the age of two, a now widely-known guideline but a difficult one to implement and a common source of worry for new parents. It may thus come as a surprise that he has criticized this blanket prohibition, and argued for a more nuanced view, that moderate ‘screen time,’ particularly with interactive media, is unlikely to pose serious problems in early childhood development (Christakis, 2014). Despite the seemingly clear lesson of his animal experimentation here, the implications he draws from it are more cautious.

Animal modeling of course has a long history within psychology, and it is perhaps more closely associated with the heyday of behaviourism than with cognitive science. Yet it has been

thoroughly altered by a discourse of mind, intention, and information-processing.¹⁶ Indeed, it is above all this guiding metaphor-system of information and computation that binds all of these disparate investigations together as a ‘cognitive science of media,’ absent any direct and well-established means of translating between the domains of behavioural, imaging, and modeling studies. What of computer modeling, though, the technique most clearly associated with cognitive science? Such studies considering media effects are very few, bordering on nonexistent. I was able to find only two, one of which uses a very simple mathematical model run to simulate the effects of various levels of spending on mass-media anti-smoking advertisements on rates of smoking cessation, based on past data (D. T. Levy & Friend, 2001).¹⁷ The second is of more interest but still fairly rudimentary, employing neural network simulations to explore the cultivation theory proposed by Gerbner and his colleagues. Their proposed effects cannot be readily categorized as ‘positive’ or ‘negative,’ but once again correlate with a perceptions of a ‘mean world.’ The researcher simulated a neural network in the software MATLAB, and trained it on a set of ‘input’ questions representing scenarios (*eg.* ‘a woman walks through a park’) and outputs representing possible conclusions, some involving crime or violence, and some neutral (S. D. Bradley, 2007, p. 457). The ‘TV’ stimuli are a priori defined as a set of more violent possible scenarios, while the stimuli defined as coming from ‘direct experience’ or from ‘other people’ are a mix of less violent or neutral scenarios, and sets of ‘light’ and ‘heavy viewer’ neural networks were trained by varying the proportion coming from ‘TV’ or elsewhere.

¹⁶ For a thorough examination of the issues around animal, and particularly mouse modeling in cognitive psychology and neuroscience today, see the recent work of Nicholas Rose and Joelle Abi-Rached (Rose & Abi-Rached, 2013, ch. 3).

¹⁷ This work is allied to other research on the effects of anti-drug and harm-reduction campaigns which makes heavy use of cognitivist language in more informal conceptual models rather than formal computational ones, *eg.* (Harrington, Lane, Donohew, & Zimmerman, 2006), which argues for an ‘activation model of information exposure’ differentiating high and low ‘cognition value’ messages and individuals with high and low ‘need for cognition,’ and a need for optimal matching between messages and audiences for persuasive effect.

As the researchers predicted—indeed, as anyone could have predicted given a basic understanding of neural networks and of the training sets supplied here—the model demonstrated cultivation effects, with a greater activation on output nodes representing the ‘TV’ answer among ‘heavy viewer’ models, and on nodes representing the ‘real world’ answers among the ‘light viewers’ (S. D. Bradley, 2007, p. 462). The researcher himself recognizes this as merely an “existence proof,” which will “not come as a surprise to connectionists” (*ibid.*). Nonetheless, he concludes that it indicates how TV can bias understandings of social reality, and that his results match well with research on human subjects. Modeling allows for much better control over stimuli, and more conclusive evidence for causation, as in many domains of science; presumably, as techniques develop for simulating more and more realistic computational brains, as discussed in the previous section, these will prove a useful supplement to neuroimaging in many domains, perhaps including that of media effects research. Yet the weaknesses of Bradley’s neural-network study indicate why this may be difficult: if one is investigating the effects of technologically-mediated messages using a computer model, how can a contrasting ‘unmediated’ message be provided, except by arbitrary definition? I return to an analogous dilemma in the context of imaging research below, which may account for the rather limited number of studies in that domain explicitly considering media. Perhaps more imaging studies than we expect are really testing media effects—because how else are you supposed to present stimuli to a subject immobilized within a brain scanner?

First I turn to some of the results which have been recruited to suggest, contrary to those discussed above, that new media are having more positive effects on our minds and brains. One of the first major studies to consider the psychological effects of Internet use, widely cited in the following years, was a longitudinal survey carried out at Carnegie Mellon which gave 93 Pittsburgh families computers and access to the World Wide Web for the first time, then evaluated their

psychological well-being using various measures over the following two years (Kraut et al., 1998). Their first finding was what they called an ‘Internet paradox:’ that although the medium was a “social technology” intended to enhance communication, and did increase total volumes of communication, individuals who began using it tended to exhibit “declines in social involvement and the psychological well-being that goes with social involvement,” feeling more lonely, depressed, and detached from others, which the researchers hypothesized could be due to the substitution of “weak ties for strong ones” (Kraut et al., 1998, p. 1029). When the researchers revisited this question a few years later with a different sample, however, they found that these results had disappeared overall, and that the only negative effect which persisted was a moderate increase in reported stress, with a contrasting increase in general positive affect and social involvement. Though far from definitive, and perhaps owing to some independent difference between the two samples, the researchers concluded that this new study (also conducted with a sample of new Internet users) suggested that some important characteristics of the medium itself had changed, leading to different effects. Perhaps there were simply more people worth talking to online, with more family or pre-existing friends now having access themselves (Kraut et al., 2002, p. 68). The second study also considered the relationship between introversion/extroversion and Internet use, finding that rather than ‘social compensation,’ where the introverted might find greater benefits, they observed a ‘rich get richer’ dynamic, where extroverted individuals with broader pre-existing social networks found greater positive effects on self-esteem and social connectedness (*ibid.*, p.58, 67). Even these broadly positive findings are thus fairly limited and nuanced: digital technologies seem to have some cognitive effects, to be sure, and they may in certain cases at certain times appear to be positive, but they are far from deterministic, seeming to change over time and in relation to different individuals’ patterns

of use. Again, this is typically recognized by the researchers themselves, but such caveats do not always accompany these findings as they are disseminated in the public sphere.

Later meta-analyses suggested instead that the overall effects of the Internet on psychological well-being were slightly negative (Huang, 2010). Given the wide variability in effects, in research methods, and the dramatic changes in the nature of Internet use across the time period covered by these studies, however, such aggregation is of somewhat dubious utility. The Internet has gone from a largely textual medium with which one engages for a relatively limited time each day—think of the now-archaic phrase “I’m dialing in to my ISP so I can check my e-mail”—to a ubiquitous hybrid of all prior media. Now many of us in the developed world have a telephone that we rarely use for voice calls, and more for its ability to ‘push’ e-mail to us continually as it is received. Digital downloads and streaming video are well on their way to supplanting the traditional avenues for accessing television and film content, while the phrase ‘social network’ conjures up not a system of real-world ties but a very specific kind of Web site whose creators have become the new face of speculative, entrepreneurial capitalism. If media are both defining of and defined by their patterns of usage, then it is hard to suppose that the Internet of 1998 is ‘the same’ medium as that of 2002, 2010, today, or ten years from now.

Video games, at least, would seem to offer a more stable pattern of interaction over time: recall that studies of their violence did not show substantial differences in effects based on the increased visual realism of more recent games. Apart from the potential association between violent content and aggressive cognition, some research has suggested more positive effects from games. One study used a ‘Flanker test’ to measure ‘attentional capacity:’ subjects were asked to indicate whether a square or diamond appeared within one of six rings on a display, while a ‘distractor shape’ appeared outside the rings. The speed of subjects’ response under different kinds of distraction

conditions was taken to indicate their overall attentional capacity or ‘resources,’ and it had previously been observed that as the tasks become more difficult, the distractor has less of an effect, seeming to imply that ‘spare’ attentional capacity ‘spills over’ more when the task is easy, but less as the difficulty increases (Green & Bavelier, 2003, p. 534). Testing a group of action video game players against non-video game players, the distractor effect was more stable in the former group, indicating that they had more free attentional capacity. Three other experiments were carried out along similar lines, testing subjects’ ability to enumerate squares briefly flashed on a screen, to pick out the spatial location of items displayed, and to perform an ‘attentional blink’ task, detecting two target symbols displayed in succession and reporting whether they were present in each run.

For each experiment, the researchers found results consistent, on their interpretation, with increased attentional and perceptual capacity amongst those who regularly played video games (*ibid.*, p.535). To further support the interpretation that this was specifically due to action video gaming experience, and not some predisposition of those with greater attentional capacity toward more gaming, they also carried out these tests after ‘training’ a set of non-gamers either on *Medal of Honor*, a ‘first-person shooter’ action game, or *Tetris*, a pattern-matching game that “demands focus on one object at a time” rather than the spatially distributed visual attention required in action games (*ibid.*, p.536). Again, they concluded that action games had a specific enhancement effect on all their measures of visual attention. Elsewhere, however, the same researchers wisely cautioned that “while a strictly dichotomous classification into ‘good’ and ‘bad’ makes for nice headlines... such a scheme ignores the fact that human experience is multidimensional; almost all experiences are ‘good’ in some ways and ‘bad’ in others” (Bavelier, Green, & Dye, 2010, p. 693). The changes in visual attention which their specific experimental apparatus picks out as ‘enhancements’ might readily be

interpreted in a different light when considering facility with media requiring ‘deeper,’ more focused and regular visual attention, such as printed books.

Despite the ambiguities surrounding them, these findings of potential positive media effects have given rise to a whole industry. The companies offering ‘brain training’ claim to have their basis in neuroscience, providing a scientifically sound means to ‘improve memory and focus’—amongst other quantifiable and visualizable benefits using smartphone apps—while requiring only a mildly diverting experience much akin to playing a conventional video game. The supposed ‘science’ is that

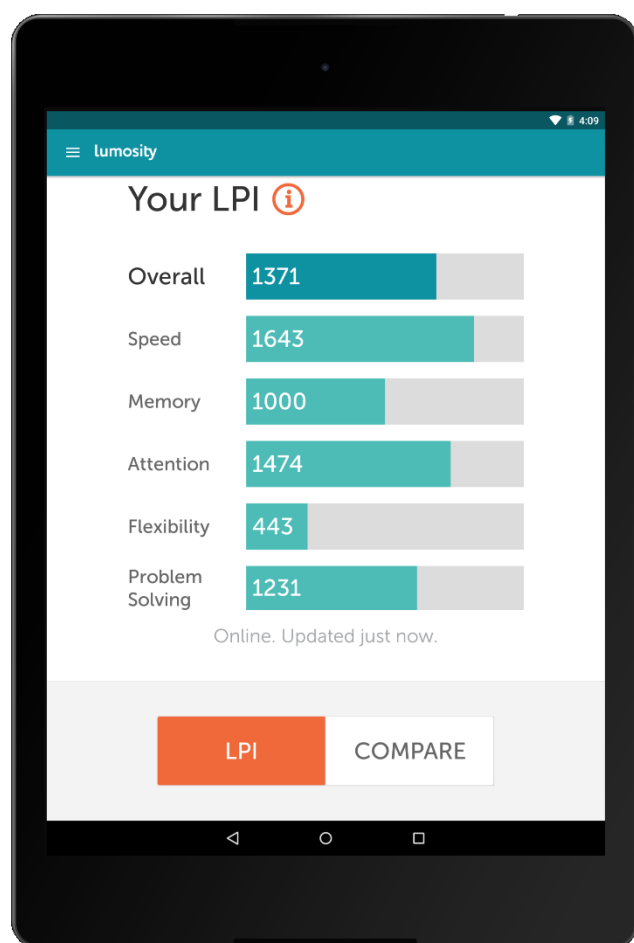


Figure 29 - Lumosity performance tracking screenshot

discussed above, which implies an improvement in attentional and sensorimotor capacities, and concomitant changes in the brain, following the playing of certain video games. Lumosity, the foremost corporate provider of specially designed brain training games, has received several rounds of venture capital funding, and its mobile app for Apple’s devices has been downloaded over 10 million times. Their advertisements promise to ‘activate a better you,’ with ‘brain games’ that ‘scientifically’ improve memory, focus, and relieve stress. Its corporate spokespeople love to repeat the obvious analogy: that just as an athlete will train with weights to improve their overall strength and thereby their specific task performance, playing these ‘cognitive games’

will produce trackable gains in general intelligence. More importantly, it invokes neuroplasticity, and operationalizes these supposed gains in a testing protocol that tracks improvements on their ‘brain performance index,’ taking the common experience of character progression from video games and leading subjects to map it on to their own real-world cognitive selves (Figure 29). Lumosity’s apps knit together games that start to feel like psychological tests, and psychological tests that start to feel like games.

Somewhat paradoxically, the scientific community has greeted purposely-designed and elaborately-marketed brain training products with more skepticism than the notion that good old-fashioned action games might improve our cognitive capacities. The hype around products like Lumosity’s has grown so intense that a community of interested neuroscientists has recently felt the need to sign a consensus statement urging caution about the benefits of brain training. Well over 100 scientists attached their name to this statement, a list which continues to grow as an addendum to its online publication. They take aim specifically at commercial enterprises which claimed to have products ‘designed by neuroscientists,’ and produced ‘exuberant advertising’ that made scientists ‘cringe’ at claimed “improvements in the speed and efficacy of cognitive processing and dramatic gains in ‘intelligence’” (Max Planck Institute for Human Development and Stanford Center on Longevity, 2014). Worse still, purveyors of these products played on the anxieties of older adults about age-related cognitive decline, and made unwarranted claims about their ability to stave it off. They argued that before investing in brain games, individuals should consider their opportunity cost: if such games replace otherwise sedentary, passive activity, they may be beneficial, but the best practice for cognitive and overall health is simply to “lead physically active, intellectually challenging, and socially engaged lives” (*ibid.*). Despite the common-sense nature of their final recommendations, this consensus statement—like many other examples of this form—serves as a clear instance of

scientific boundary-work, an effort to demarcate the community of serious, objective scientific researchers examining changes in cognitive function, from the wider world of quackery and hucksterism.¹⁸

Not all researchers are willing to be recruited into this supposed consensus, however. Further muddying the waters, another group of 100-plus scientists recently produced a new open letter critiquing the first, and arguing that it “did not reflect a true consensus from the community” (“Cognitive Training Data,” 2015). While also distancing themselves from claims that particular products were supported by scientific evidence, and calling for better peer-reviewed research into their efficacy, they took issue with the claim in the first statement that there was simply ‘no compelling scientific evidence’ for the usefulness of computer-mediated brain exercises in reducing or preventing cognitive decline. Their letter argued, on the contrary, that there was a growing body of scientific research showing that specific types of cognitive training could drive specific types of improvements, and that the original consensus statement ignored the powerful and continuing role of neuroplasticity throughout the lifespan. By failing to recognize the results of these well-run studies, the statement may contribute to a slowdown in funding for academic research in this area, and thus “precisely the environment that you (and we) seek to discourage – one where investments in science are outweighed by investments in advertising” (*ibid.*). This response was spearheaded by Michael Merzenich, a prominent researcher in the area of neuroplasticity, a co-inventor of the cochlear implant, and the most important scientific interlocutor in Nicholas Carr’s book. As the list of signatories to his letter duly reports, he is also in a conflict of interest position, having co-founded three companies offering brain training products. The overwhelming majority of signatories have no such conflicts, however. The difference of opinion seems not to be a matter of financial investment

¹⁸ Several of its signatories are researchers already cited in this chapter (*e.g.* Daphne Bavelier and Shawn Green).

but disciplinary affiliation: the majority of signatories to the initial statement were behavioural researchers, focused on examining claims for cognitive training through aptitude tests and other psychological instruments, while the majority of those signing the response were neuroscientists, focused on examining claims for brain change through neurophysiological techniques (Horvath, 2015).

Corresponding with the outward-facing boundary work, then, of distinguishing true science from marketing-driven overclaims, there is a policing of boundaries within scientific communities. Particularly now that they have moved online, the *genre* of the scientific consensus statement seems less definitive than ever in achieving interpretive closure. While one may continue to garner hundreds more digital signatories post-publication, it may quite readily and rapidly lead to the mobilization of a counter-consensus, often deriving from different disciplinary backgrounds, and focused on different research methods and goals. At stake here again are questions of media effects: how does playing games change the brain, and can we design electronic games to produce positive effects on cognition? Beyond the typical issues of underdetermination in science, the shifting frontiers of technology add a further layer of uncertainty, with new developments in game design and in brain imaging making consensus ever more elusive. Even some of those who signed the first statement criticizing commercial brain-training products and their marketing are in the process of developing their own brain-training software with commercial aims.¹⁹ Hence this boundary-work sometimes works to occlude both a considerable internal heterogeneity, and some broader points of consensus. Foremost among the latter is the view which first became mainstream within early

¹⁹ Adam Gazzaley, whose product *Neuroracer* is in the process of FDA approval, is one researcher whose unusual position as both a critic and promoter of brain training has received some attention in the popular media, with headlines asking ‘Can Video Games Fend off Mental Decline?’ (Thompson, 2014). Gazzaley at times sounds more in line with the ‘counter-consensus,’ having pushed the first group to use “less pessimistic rhetoric,” also fearing that “excessively negative statements might scare off research-funding agencies” (*ibid.*).

cognitive science: that the mind and brain are to be understood as information-processing systems. From the early days of ‘limited effects’ theories in media studies, we have arrived at conceptions of mediated effects channeled through ‘limited-capacity’ cognitive processing (Lang, 2000). More specifically, though, there is a building consensus that games are another important class of information-processing system that seem to reshape ours in the playing. We of course become more proficient in the games themselves if we play them for an extended period of time, as with any skilful activity, and this is of necessity accompanied by changes in the brain. Still up for debate, though, is the scope and transferability of these changes, along with whether particular games can be engineered to produce stronger effects.

More fundamentally, the appeal of games—particularly, but not exclusively electronic ones—comes from their capacity to sustain our attention. Though an ‘acquired taste,’ and one involving a seeming generational divide, for the initiated, there are few experiences so readily able to produce a ‘flow’ state of rapt concentration and focused attention as video games. Other intellectual or physical pursuits can certainly induce this state as well, but by no means as easily. At the root of this effect, it seems, is the dynamic adjustment of difficulty as user proficiency increases: a ‘balancing of challenge and skill,’ which even in a stripped-down mathematical task can induce flow states, but is most strongly associated with video games (René Weber, Tamborini, Westcott-Baker, & Kantor, 2009; Anguera et al., 2013; Ulrich, Keller, Hoenig, Waller, & Grön, 2014). fMRI research has used this principle, both in specially designed computer-delivered tasks, and in commercially available video games, to study the flow state; while conclusions are highly speculative as in most basic imaging research, some have persuasively suggested that flow must involve a synchronization of attentional and reward networks within the brain (René Weber et al., 2009). While in a sense this is simply translating the intuitive concepts into a neuroscientific frame, computational imaging

techniques promise to further elucidate the interconnected systems of neurons which make up these networks, through forms of pattern-matching which move beyond localizing specific phenomena in ‘brain areas for *x*.’ Claims for general-purpose brain training games seem to be quite dubious, by a wide consensus of serious researchers; one might as well play *Call of Duty*, or far better do some strenuous physical activity. Yet the efficacy of games for inducing flow, and the evidence for their positive effects on cognition, means they retain their appeal for many of the same researchers, particularly in staving off age-related cognitive decline, and treating patients in rehabilitation, for instance from strokes or traumatic brain injury. In these contexts, concerns that the games may be displacing other valuable sorts of activity seem less apt.

Yet in the wider debate over gaming, this seems to be at the root of many concerns. A game which promotes shifting attention and multi-tasking targeted toward seniors may have demonstrable benefits: ‘Neuroracer,’ a game of this type, is presently undergoing review as a medical device by the FDA in the US. Yet pundits, parents, and many others worry about the implications for coming generations of having been raised on exponentially more complicated games. As one online commenter quipped, “it’s only fair, that Video Games would fend off mental decline in the elderly, since they seem to bring it on in the young” (see Thompson, 2014). The games, moreover, are merely one strong competitor among many for the attention of youths. The multitude of ‘screen-based media’ available seem to be exerting ever greater pull, outcompeting more traditional pastimes such as organized sports and musical instrument-playing. This speaks again to the ‘generational divide in cognitive modes’ proposed by Katherine Hayles and extended by Bernard Stiegler, between ‘hyperattention’ and ‘deep attention.’ The paradigmatic figure for the latter is the scholar enmeshed for hours in a thick book, blocking out all outside stimuli in a carefully cultivated flow state; the figure of the former is the student hopping back and forth between myriad easy sources of flow,

pouring a half hour here into following a chain of online scholarly references, there into playing one video game, some into watching professionals play another, all while partially attending on a continuous basis to pop-up email notifications, household tasks, and unbidden compulsions to check frequently-updated websites. As one who often resembles the second figure more than the first, with Hayles I encourage an openness to the powers, pleasures, and perils of both cognitive modes. This nuance is often lost in popular debates, which tend to pit traditionalism and moral panic against unbridled techno-optimism: either games and mobile phones are destroying our children's sensitive brains, or 'everything bad is good for you' (S. Johnson, 2005), and we can remake ourselves into brilliant multitaskers through video games. Equally often lost is a healthy skepticism in the interpretation of neuroimaging results.²⁰ Though valuable exceptions can be found, the mass media is often a poor interpreter for the science of media effects.

From my review, I conclude that this scientific literature is ultimately rather sparse, and despite the strong conclusions some wish to draw, many open questions and controversies remain. This should not be taken to imply, however, that we know nothing about the effects of media on the brain and mind, or that these are somehow irreducibly complex. Behind many of the controversies discussed lies a rough consensus on more fundamental issues. We may conclude definitively that digital media have some influences on our brain, particularly on the brain systems responsible for attention and executive control. This finding has been 'triangulated' by various imaging technologies and behavioural methods, and so it seems solid despite various associated uncertainties. Far more research is needed to determine the scope of these effects, not to mention simply understanding the neural networks underpinning the basic phenomena of attention. Yet rather than viewing new media

²⁰ As we shall see in the specific case of ADHD debates, these tendencies are especially pronounced in the British tabloid press, and in populist-conservative leaning American papers like *USA Today* and the *Washington Times*.

forms as radically distinct in their capacity to produce ‘dopamine squirts’ or ‘hijack’ our brain systems, it seems we are better served by regarding them as highly effective producers of flow states. Their use and abuse are best understood not through the framing of ‘addiction,’ but as phenomena which outcompete other activities within an economy of attention. There is no reason not to enjoy these ‘easy’ producers of flow, or to forbid children from enjoying them—be it playing a game or idling away on the Internet—provided that they do not displace harder, but more rewarding pursuits requiring ‘deep attention.’

Precisely how much displacement we as individuals, families, and societies should allow remains a matter for public debate and personal decision-making. While we may conclude with some certainty that these effects exist, and call for further research in understanding them, the implications to be drawn about the positive or negative consequences of these effects lie beyond the reach of any scientific research, and will forever remain disputed and underdetermined. One’s view depends considerably on how strongly one feels about the mode of deep attention long associated with literature and print, and how profoundly one believes it to be threatened by the generational shift toward differing patterns of attention. The best commentators are those who remain open and reflexive about this historical link between print media and public rationality, whatever side they fall on and whatever evidence they cite. The worst are those who suppose that digital media are some wholly new force befalling the normal, natural human brain purely for the worse, or those who adopt a starkly instrumentalist view, believing that human rationality as extensively supplemented by information technologies will be just more of the same, only faster and better. As espoused and actively pursued by such luminaries as Google’s Larry Page,²¹ this is in a sense the ‘limited effects’

²¹ As Page once put it, “Ultimately I view Google as a way to augment your brain with the knowledge of the world. Right now you go into your computer and type a phrase, but you can imagine that it could be easier in the future, that

model of our time, and it is as troubling as the alarmism spouted by the likes of Susan Greenfield. Whatever technologically extended rationality may be, it is clearly on its way to becoming something *different*. Another prominent way in which this is occurring has been popularly labelled the ‘Google effects on memory:’ some research suggests that individuals proficient with the search engine tend to treat it as a ‘transactive memory’ system, becoming less likely to recall specific pieces of information than how to recall them (Sparrow, Liu, & Wegner, 2011). Page dreams of one day affording a direct brain interface to this brand of transactive memory, and beyond that of linking human minds up with artificial ones in an even more profound symbiosis. At least for now, though, while changes to our brains are necessarily part of the story of media effects, they are by no means the whole story, as some commentators seem to suggest. If media are changing our brains, it can only be by way of changes in our patterns of attention and interaction, both with technologies and with each other.

In closing this section, I return to cultivation theory, and outline a potential synthesis with cognitive science. I contend, following McLuhan and Andy Clark, that there is nothing more natural to human thought than extending our capacities through technological media. Through both form and content these technologies then come to shape our view of the world. As L. J. Shrum has suggested, perhaps the most important mechanism underpinning the sorts of effects described by Gerbner and the Cultural Indicators Project is heuristic processing (Shrum, 2008; Shrum & O’guinn, 1993). Heuristics are in turn integral to cognitive science, as I have discussed in previous chapters, both in the programs it developed and in the general sense that the field was guided by a tools-to-theories heuristic, conceiving the nature of the mind through digital computing. Gerd Gigerenzer,

you can have just devices you talk into, or you can have computers that pay attention to what’s going on around them... Eventually you’ll have the implant, where if you think about a fact, it will just tell you the answer” (S. Levy, 2011).

who proposed this particular heuristic, is one of the leading contemporary scholars of heuristics more generally and of ‘bounded rationality.’ That approach was first proposed by Herbert Simon himself as an alternative to classical utility-maximizing conceptions of human nature—*Homo Economicus*, as the construct is often labelled—and as its name suggests, it considers cognitive agents to be graced with more limited capacities. These constraints are of various kinds: temporal, physiological, affective, and so on. In practice they imply that we are finite beings who can never access or evaluate *every* relevant piece of information in forming a judgment or reaching a decision. While we do often strive to deliberate and weigh many competing considerations in a conscious, classically rational fashion, more often we rely to some degree on ‘fast and frugal’ rules of thumb (Gigerenzer, 2008; Goldstein & Gigerenzer, 2002). These supply a crucial vector for media effects, inasmuch as the base of knowledge on which these heuristics operate is drawn as much—if not more—from mass media as from personal experience.

The same could be said of the raw materials for more conscious reasoning, but the consequences of media for heuristic processing are more unique and worrisome in that they operate below the level of full awareness. In other words, they cultivate certain ‘naturalized’ ways of seeing the world. Two of the most important heuristics, for instance, are those of *availability* and *recognition*. We often rely on how readily some information comes to mind as a proxy for some other feature of interest. Mediated messages, in turn, have a powerful influence on what is available or recognizable to our cognition. The former heuristic is the reason that we overestimate the frequency of plane crashes, terrorist attacks, and violent crimes generally, as opposed to other risks: because they are more widely covered in the media. The latter heuristic underpins the very idea of ‘branding,’ in that purchasing decisions are more commonly made based on a sense of familiarity rather than an objective assessment of selling points. As with other areas covered in this section, further research is

sorely needed to explore this intersection of bounded rationality and communications technology, bringing it into dialogue with neurophysiological accounts. We must likewise be attentive to ways that commercial and government actors may employ knowledge of heuristics to deliberately cultivate certain mentalities. Heuristics are useful when, say, you're asked the capital of a state and you only know the names of a handful of cities there—the first one that comes to mind is likely a good guess, though this strategy is certainly fallible. They may have negative though inconsequential results when you're deciding where to eat dinner, leading you toward some mediocre chain instead of a local gem. More troubling consequences ensure, however, when large masses of people rely chiefly on such heuristics in choosing how to cast their votes, or when demagogues leverage the view of a 'mean world' to promote repressive policies. Boundedness and heuristics are not to be seen as defects afflicting us in comparison with some hypothetically ideal cognitive agents; such agents are fictitious, and as Gigerenzer has suggested, we often do better to rely on these rules of thumb rather than trying to recall and weigh every relevant bit of knowledge. External mediations help us to transcend many of our cognitive limits, and without the archival and transmission of knowledge through technology, reason as we understand it could not exist. Yet by the same token we must recognize when media are cultivating unwanted habits of thought, and when we need to employ slower, more deliberative and explicit cognitive processing.²²

²² Adopting terminology from the psychologists Keith Stanovich and Richard West, Daniel Kahneman has recently written a book examining 'System 1' and 'System 2,' heuristic and deliberative reasoning, as distinct but coexisting streams of cognitive processing. While certainly not the first to suggest such a division, his argument is compelling, and he now presents a wider view of System 1 than as a source of error as compared with classical idealizations (Kahneman, 2011; *cf.* Petty & Cacioppo, 1986; Sloman, 1996). Effective human reasoning requires both, and the question of which is 'better' is a meaningless one without asking 'what for?'

Another angle on mediation in science and cognition.

Despite notable public concern with this issue, I found relatively little direct engagement with the question of media effects from a cognitive science perspective. In another sense, however, digital media are everywhere in cognitive science. Computing technology, as discussed in the last section, was constitutive of the discipline in many ways. The paradigm of interactive computing advocated by McCarthy, Newell, Licklider and others within the early history of cognitive science has spread out across the sciences, and from laboratories into our homes and everyday lives. The domain in which the role of computation has received perhaps the most attention is that of brain imaging, with some STS scholars having sought to unpack the mediated processes of construction behind fMRI and PET images (Beaulieu, 2002; Dumit, 2003), and myriad other commentators in recent years offering their own wider critiques of these methods and their ‘seductive appeal’ as tokens of persuasion (Weisberg, Keil, Goodstein, Rawson, & Gray, 2007; Uttal, 2011; S. L. Satel, 2013). These critiques all note that these scanning techniques observe blood flow in the brain, which serves at best as a proxy for the average firing rate of many neurons within ‘voxels,’ and that many steps of standardization, selection, and correction are necessary between the raw scanner data and the persuasive colour images of brains. Hence as most imaging researchers readily admit, they should not be taken as unmediated windows into the true nature of the brain, but as products of a difficult and fallible construction process whose methods should be open to critique.

Throughout this process, computational tools do not supplant but complement more traditional modes of extending human cognition such as oral communication and gesture (Alac, 2008). Some of the techniques of ‘machine learning,’ one of the more recent names for the heuristic programming approach once called AI, seem to promise ways beyond the localizationist approach of simply highlighting standardized brain areas with the greatest activation, instead helping to uncover

distributed networks and patterns in the brain associated with particular mental states. Researchers can feed a subset of imaging data which varies along certain dimensions—for instance brain scans of subjects with and without a given mental disorder—into a machine learning algorithm, ‘training’ it to discriminate by feeding it the correct classifications, and then testing its accuracy in classifying the full set. This allows for common patterns to be picked up across scans which may be imperceptible to human operators; it also allows for new ways of potentially overstating significance and engaging in questionable research practices.²³ Intuitively, if a classifying program succeeds in distinguishing between brain scans correctly better than chance—*ie*, more than 50% of the time for a 2-class problem, or 25% of the time for 4—then it has hit upon some underlying correlation between brain states and a phenomenon of interest. As a recent paper contends, however, this approach has a limitation that most in the machine learning community understand, but those applying its techniques in neuroscience often do not. The theoretical threshold only applies to infinite data sets. For small data sets, these algorithms can often perform substantially better than chance at discriminating between random data sets without real differences: for sets of 20, they can sort them ‘correctly’ as much as 70% of the time. Thus neuroimaging studies having that small a population, as they often do, ought to surpass this threshold rather than the theoretical one (Combrisson & Jerbi, 2015).

Taking for example one frequently-cited paper which claimed to ‘read hidden intentions’ from brain scan data using machine learning techniques, they scanned three male and five female subjects multiple times while asking them to choose between one of two possible tasks. They then trained a ‘linear support vector pattern classifier’ on seven of these eight data sets using the correct

²³ This phrase has become such a focus of concern within the neuroscience community online that it has become a widely-known three-letter acronym: QRPs.

information about which task the subject selected. By testing different clusters of voxels within the eighth data set to see if their classifier could accurately discriminate between the tasks the subject had intended to perform, in advance of their action, they achieved at best a 71% accuracy when it was applied to specific regions of the prefrontal cortex. From this they concluded both that they could read intentions from scan data, and that ‘intention-related information’ was encoded in these particular regions of the cortex (Haynes et al., 2007). Yet while they emphasize the strength of their result in comparison to the theoretical chance threshold of 50%, in fact they may have barely exceeded the chance threshold needed to achieve a truly significant result using their machine learning approach.²⁴

As these new techniques allow researchers to analyze brain scans in new ways, they equally confer the freedom to screw up or fudge data in new ways. Unfortunately, nuance and critical reflexivity are often elided as these images are disseminated through popular media. A great deal more could be said of neuroimaging research and the specialized computational tools by which it is produced, but this lies beyond the scope of my project here. In concluding this section, I wish to call attention instead to the ways in which more everyday media technologies are coming to transform research practices in the cognitive sciences. As one researcher recently observed, new modes of “post-publication peer review—whether instantaneous and occasionally flippant as on Twitter, or more nuanced commentaries as on blogs, and of course on dedicated websites like PubPeer” mean that while researchers who run ‘underpowered’ studies or ‘double-dip’ for their statistics might make it through the fallible process of standard peer review, the world will learn of their transgressions

²⁴ Their study used a total of 8 subjects and 32 trials, making their significance threshold for $p < 0.05$ somewhere between 62.5 and 70% classification accuracy, depending how one interprets the results of Combrisson and Jerbi, and Haynes et al.’s use of multiple same-subject scans; for more stringent p -values their results would not pass the threshold (Combrisson & Jerbi, 2015).

“within days after it appears online,” or “within hours” if they accompany their “precious new study with an embellished press release” (Reas, 2015). Writing pseudonymously online as ‘practiCal fMRI,’ this researcher is one of many who have brought into the open once-hidden networks for informal circulation of critiques and disputes both polite and snarky, along with knowledge about the ‘nuts and bolts’ (their phrase) of experimental apparatus. Discussions carried out through such online channels have produced, for instance, novel techniques for using low field strength MRI to diagnose mild brain trauma, which is poorly captured by traditional scanning methods, as well as ‘crowdsourced’ lists of potential physiological and pharmaceutical confounding factors for blood oxygenation in the brain (and by extension fMRI results).

Another sense in which the Internet has entered into psychological research in a transformative way is through the use of Amazon’s Mechanical Turk service, an online platform branded as ‘Artificial Artificial Intelligence.’ It connects task creators, known as Requesters, with thousands of distributed online human workers, to carry out ‘Human Intelligence Tasks,’ or ‘HITs,’ rapidly for tiny sums of money.²⁵ These are often such prosaic tasks as the completion of ‘captcha’ forms, themselves online agents delegated the task of sorting out human from nonhuman intelligence. More recently, however, psychologists and social scientists have begun employing the service as a sample population for completing surveys and experimental tasks. One study evaluating

²⁵ Typically a worker receives around a dollar an hour, or as one group of social science researchers commenting on the use of the service put it, “nickels and dimes for 5-minute tasks” (Buhrmester, Kwang, & Gosling, 2011, p. 3). Originally the service could only pay cash to users within the United States, requiring that others around the world accept Amazon gift cards, but this requirement has recently been lifted. Consequently the share of task completers in India and elsewhere in Asia has increased substantially, where the purchasing power differential makes these low US dollar wages comparatively more appealing. Amazon’s tools for task creators include options for filtering out workers based on nationality, language proficiency, or test questions, facilitating cross-cultural comparisons in the context of research. As others commenting on the usefulness of Mechanical Turk have observed, it may be preferable both ethically and in terms of acquiring good research subjects to pay wages that are high by comparison to those typically offered on Mechanical Turk, and closer to those offered to participants in laboratory settings (Crump, McDonnell, & Gureckis, 2013).

the uses of the service for cognitive psychology employed classic experimental tasks like the Stroop color-identifying task and the Flanker letter-identifying task, for instance, and concluded that nearly all the response patterns obtained matched those “established in more controlled laboratory settings” (Crump, McDonnell, & Gureckis, 2013). In general, while samples from Mechanical Turk are far from ideally random—being necessarily composed of those willing to perform simple tasks online for minimal pay—they compare favourably with those conventionally used in psychology, particularly undergraduate students. When the alternative is also not particularly representative, Amazon’s service offers researchers an appealing way to gather large, relatively diverse samples and simplify the administrative task of paying participants.

Other mass-market online platforms have also been used in recent years to collect psychological data. One notable experiment explored the phenomenon of ‘emotional contagion’ using the social networking site Facebook, with the company’s blessing and assistance (A. D. I. Kramer, Guillory, & Hancock, 2014). This analysis showed the potential power of such data-gathering techniques, in the scope of its sample (689,003 users) and in the nature of its conclusions, which showed that when the frequency of positive or negative terms in users’ “News Feeds” on the site was skewed, it produced a corresponding increase in expressed emotional positivity or negativity *by* that user in the following week. The online response, however, was less one of approval with regard to the significance or novelty of these results, than of concern with the ethical ramifications of such an experiment. As one popular account phrased it, the study sought to answer “the question of whether emotional states can be transmitted across a social network. Result: They can! Which is great news for Facebook data scientists hoping to prove a point about modern psychology. It’s less great for the people having their emotions secretly manipulated” (W. Hughes, 2014). Though the experiment had passed some level of review by the two academic authors’ own Institutional Review

Boards—the study’s lead author, Adam Kramer, is a Facebook employee—the response after publication indicated that they had failed to consider the real implications of how users’ emotions were being manipulated without consent, particularly given the positive results of the experiment.²⁶

It also served as another indication of how post-publication review is transforming: ethical concerns were rapidly circulated on specialist blogs, and then spread into the mainstream press following an outcry by users (Albergotti, 2014). The paper itself was also open-access and included a comments section of its own, which quickly became active with other researchers and the wider public voicing their concerns. The journal helpfully offers a comprehensive record of the article’s online impact, noting among other metrics that it was picked up by 173 news outlets, 3779 Twitter users, and 412 pages on Facebook itself, giving it as of March 2015 the highest ever online impact figure for that journal (*Proceedings of the National Academy of Sciences*). Widely used social networks like Facebook are thus potentially valuable sources of data, but must be approached with caution. New possibilities for digitally mediated data collection bring with them a new scope of potential harm and a duty of ethical care, while the very popularity which makes these sites useful potentially amplifies the public scandal should an ethical breach occur. Digital publication likewise brings with it new possibilities for quick turnaround of research, and for quick responses both positive and negative, sometimes coming from unexpected quarters.²⁷ The furor was ultimately sufficient that Facebook’s chief operating officer, Sheryl Sandberg, felt compelled to issue a public apology for the ‘poorly

²⁶ As one journalist recounts, “Cornell University’s IRB had taken a look, but concluded that the study did not require full review. Why? Because it was Facebook’s team alone that carried out the experimental manipulation of news feeds and collected the results before handing them over to Cornell’s main researcher” (Hsu, 2015). In retrospect this IRB position seems to raise even broader concerns about ongoing academic-industry collaboration in this sphere than would a simple mistaken decision about this particular experiment.

²⁷ Some of the negative implications drawn by assorted commenters on the original paper cover not only the ethics but the substantive content of the study, noting that the scandal may distract from the small effect size and other methodological concerns, while Facebook may well be leveraging the publication of research showing it can influence users’ emotional states in order to raise its advertising rates. *PNAS*, for its part, is happy to aggregate all of this outrage into a single ‘Online Impact’ metric which brackets all content and makes the article look like a real blockbuster.

communicated' fashion in which this research was carried out, and promised more institutional review of research by a panel within the company. Her phrasing made clear, however, that this kind of research on user responses was 'ongoing,' leading some commentators to conclude that Facebook was at best 'almost' sorry (Rushe, 2014).

Another instance in which a digital media company has converted its user base into the subjects for a 'virtual psychology laboratory' was more positively received, and earned favourable comparisons to the Facebook study (Hsu, 2015). Riot Games operates the wildly successful *League of Legends* game, which is presently played online by an estimated 27 million people daily. The game itself is free to play; the company makes its money primarily by selling virtual characters and items to players—especially those who play a lot—and secondarily by operating and selling tickets for a league of professional players with tournaments played in front of live audiences. The competitive nature of the game means that some players are prone to extremely abusive communication through in-game messaging functions. Here as elsewhere online, this is known as 'trolling,' and it can easily drive players away. Hence the company's in-house research team has devoted considerable effort to understanding the psychology of players who engage in trolling, and ways of reducing such behaviour in-game. They have largely maintained the goodwill of their userbase by transparently communicating their objectives—widely shared by the majority of players, who rate such trolling among their least favourite parts of the game—and by enrolling the community in policing mechanisms such as a 'tribunal' which decides punishment for repeat offenders.

Riot and countless other companies operating online routinely engage in the same form of experimentation as Facebook did in its study. The practice is known as 'A/B testing:' randomly divide a sample of your customers into two groups, and instead of delivering your content in the standard way, you feed them distinct new versions and see how the differences affect user

behaviour. As ‘data scientists’ with Riot put it, they began by working from longstanding psychological theories, like colour priming: testing whether differently coloured messages would impact player behaviour differently. They found that, among Western gamers, a red message “warning about the counterproductive results of negative behavior—such as, ‘Teammates perform worse if you harass them after a mistake,’” was more effective than a ‘control’ white message in reducing trolling behaviour, while blue messages highlighting the benefits of politeness had the same effect. Conversely, they could also tell that these colour associations were different amongst the large population of Chinese players, and so the same priming effect did not hold (*ibid.*).

Eventually, however, they began designing more novel ways of manipulating player behaviour. Among the most effective of these is a ‘limited messaging’ feature used to punish players who have been reported multiple times for insulting others. Players thusly punished are limited to 5 chat messages in each phase of the game, and banned from sending messages in ‘all chat,’ visible to the other team; by randomly assigning some players to this style of punishment and others to a ‘control,’ a total ban on messaging, Riot can easily conduct a large-scale virtual experiment, with rapid feedback as to which group was less likely to reoffend, and which groups reported greater overall enjoyment. Much like Facebook’s experimentation with users’ emotional states, this is a form of A/B testing widely used by corporations operating online: the Twitter home page, which leads directly to signup for new users rather than providing any glimpse of the content, is one site that has been remade through extensive A/B testing; Netflix, likewise, is continually tweaking the way it presents content to users to determine what is most likely to keep you watching (and subscribing). Riot, however, is somewhat unique in its transparency, and in the goodwill it has garnered through publicizing its experimentation. By keeping the focus on policing toxic behaviour and improving the social aspects of the gaming experience, as well as by recruiting players for community initiatives,

they have avoided the kind of backlash that Facebook experienced when it turned itself into a ‘virtual psychology laboratory.’ Their experimentation was equally oriented toward manipulating users’ emotional states, but in a positive direction, and with full openness on the part of the research team. Yet we can be sure that Riot, like other enterprises, is continually engaged in a less savoury form of A/B testing. Even its ‘pro-social’ research is ultimately in the company’s self-interest, since happier players are apt to spend more money. Yet *League of Legends* is a free-to-play game, which makes money only by players spending money on virtual characters and items. Hence a far more important, but rather less publicized side of its A/B testing is likely oriented toward maximizing player spending, just as Facebook and Google are continually testing different layouts to drive users to click on advertisements more frequently.

Pioneers of the ‘virtual laboratory’ approach to psychology, like Brian Nosek, have expressed both enthusiasm about the possibilities for industry collaboration, and concern about the status of the data thus collected. As he put it, if those databases were opened to academic researchers, “research in human behavior would advance very rapidly and change the character of how research could be done... You could imagine with this sort of iterative process that science would just come out, boom, boom, boom” (Hsu, 2015). Quite apart from the sociological issues one might raise against this iterative, factorylike model of scientific production, however, there are signs that industry is not keen to fully open its data sets to independent researchers. Riot is perhaps leading the way in fostering research collaborations with universities, with several ongoing at present, but as with the Facebook study these are designed and sometimes led by in-house research teams, providing controlled access to particular subsets of the data. Researchers cannot simply delve into the full records to unearth their own correlations, unsupervised by corporate employees tasked with protecting the reputation of their firm.

This makes for a crucial distinction between these resources and traditional sources of archival and experimental data. While further collaboration between digital media companies and academic researchers is inevitable, and will likely make it easier in the future to run

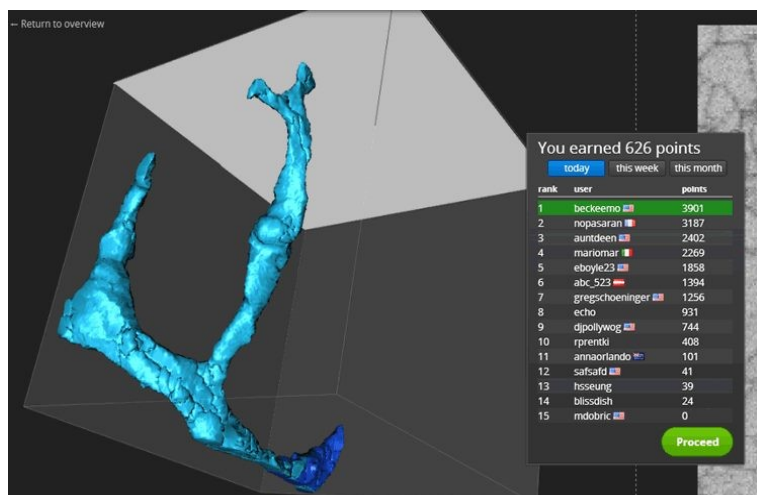


Figure 30 - Eyewire screenshot

experiments with large sample sizes, it raises significant new ethical questions. What data *aren't* being shown to independent researchers? What uses might corporations put studies to after the fact? The example of Facebook manipulating user emotions in a negative direction, and then potentially leveraging its study to command higher advertising rates – despite its small effect size – makes for a sobering case study. At a minimum, corporate research teams engaged in collaborations with universities should ensure that their work is subject to full IRB scrutiny, something which Facebook failed to do, but Riot Games has insisted upon. Ideally, though, if these corporations seek to make real contributions to science, they should be establishing timelines for opening their full data sets to credentialed researchers in decades to come, once trade secrets and user confidentiality cease to be relevant concerns. Yet another, potentially preferable approach for creating online ‘virtual laboratories’ is for researchers themselves to ‘gamify’ some important aspect of their project. This approach was pioneered by ‘Eyewire’ (Figure 30), software which originated in and continues development in association with the aforementioned Sebastian Seung’s lab at MIT (eyewire.org). In Eyewire, the search for functional interconnections between neurons is turned into a maze-like game, where players earn points for tracing out pathways in electron-microscope scans of a human

retina. The more they can improve upon the ‘best-guess’ of a machine-learning algorithm, the more points they earn. Users are thereby situated within a heterogeneous, globe-spanning assemblage, linking up multiple laboratory sites with human and non-human cognitive agents: a game that enacts distributed cognition, all in service of improving our understanding *of* cognition.

While games are on their way to becoming virtual psychological laboratories, the technologies of gaming are likewise becoming integral components of more conventional laboratory sites. Some of this research constitutes an updated form of the same kind of gaze-tracking research pioneered by McLuhan and Hurst decades ago. Rather than passively monitoring the gaze patterns of experimental subjects, now gaze-tracking systems can be used in conjunction with virtual environments and brain scanners to produce dynamic scenarios, in which aspects of the virtual environment or simulated agents within it can change their behaviour in response to the subject’s gaze (Wilms et al., 2010). Such interactive eyetracking setups can allow subjects to control elements of a virtual scene using their gaze, they can also be used to create simulations of ‘online’ social interaction, in which subjects converse with virtual humans who respond to their gaze. This is seen as a promising route to greater insight into social cognition in more plausible scenarios than the typical laboratory situations in which social psychology is studied through “tasks involving inert social stimuli (*offline cognition*)...without actual social interaction” (Pfeiffer, Vogeley, & Schilbach, 2013). Such virtual reality setups are seen as offering a way past the “enduring tension...between ecological validity and experimental control” in psychological research (Bohil, Alicea, & Biocca, 2011). They allow for the kind of dynamic, online interactions characteristic of observing social cognition ‘in the wild,’ so to speak, while permitting the kind of thorough control and monitoring required for experimental research.

While on the one end brain scans use blood oxygenation as a proxy for neural activation, on the other, they often use some mediated presentation of a stimulus as a proxy for the real world. This has long been the case, and it is one of the reasons it remains difficult to specifically isolate the question of media effects. From photos and videos this is taking on new and more advanced forms, but this research will often be construed as less about the effects of a specific mode of presentation, and more about reactions to content. Researchers in this field would do well to consider the former as well as the latter. Nevertheless gaming technologies are ideal for exposing subjects immobilized in scanners to dynamic stimuli; this interest in gaze-tracking by psychological researchers is developing in parallel with a renewed interest in virtual reality devices by the gaming industry. Though some interactive eyetracking setups for are purpose-built for scientific research, others have been developed as add-ons to commercially-available hardware for virtual reality, such as the Oculus Rift (Geuss, 2014). Better psychological research into the cognitive states produced through interaction with virtual environments would be equally useful for understanding media effects and for designing new, immersive media experiences. Specifically, incorporating eyetracking into VR systems bears promise for reducing nausea, the most notable bodily side effect of these new technologies (*ibid.*). These prospects for advancing our understanding of mediated cognition are tied to the eminently capitalist institutions which have made digital gaming such a profitable enterprise and potent affinity group within our culture. To what extent this will lead to real enlightenment on psychological questions, or to more effective marketing and monetization, remains an open question—as it was with McLuhan’s work, which likewise fascinated businesses and executives seeking profit opportunities in the nascent attention economy.

In concluding this overview of media effects and cognition, one thing should be clear despite many questions that remain unresolved: far from being limited, as the early pioneers of

effects research suggested, the consequences of new media technologies are profound, both for our own lived experience of cognition and perception, as well as in the practices by which we come to understand these processes scientifically. Neuroplasticity means that our brains are constantly changing through our experiences, including the increasing proportion of those which are digitally mediated. This suggests that McLuhan's claims of technological 'prosthetization' did indeed have some basis in reality, as did many of his other probes. Still, the extent and normative valence of these effects remains uncertain, and they are surely not unidirectional effects operating 'without resistance,' as he sometimes claimed. Beyond the effects of new technological forms, there are significant effects on the level of content as well, as discussed in relation to cultivation research. The hodgepodge of heuristics, affective states, and deliberative processes composing our bounded rationality is shaped as much—if not more—by what we read and see through media outlets than by the face-to-face channels once considered evidently 'primary' (Pooley, 2006a). Rather than scientific research being a neutral source of insight into media effects, it is one of the most active sites in which they occur.

Digital technology is restructuring every aspect of scientific production, from the research practices within laboratories to the composition and dissemination of peer-reviewed publications. For some, this process is proceeding too slowly; for others, too quickly. Issues within scientific publishing that had long been concealed behind gatekeeping institutions are now fodder for public discussion and critique, as with Ivan Oransky's 'Embargo Watch' blog, which documents problematic cases of 'embargoed' research.²⁸ Oransky has also been a prominent publicizer of

²⁸ These are findings which have been released to the press, with the stipulation that they not be published until a specific 'embargo date.' This practice has its uses, namely in giving more time for journalists to write in-depth stories on important new findings, in advance of the institution's own announcement; it conversely poses issues in terms of 'message control' by research institutions, sometimes barring journalists from conducting interviews with dissenting researchers and acquiring a broader perspective on the embargoed study before preparing their initial article about it for publication.

retracted findings, and an advocate for more thorough post-publication review. Journals are increasingly being converted into quantitative data purporting to represent their ‘impact’ and that of individual articles within them; these metrics are in turn being coupled in dubious ways to academic hiring, promotion, and funding practices. On the other hand, researchers themselves are proposing new ways to apply the same kind of ‘crowdsourced’ reviewing that Yelp applies to restaurants, on both ends of the publishing process: with utilities like ‘Journal Selector’ which filter and rank journals on the basis of speed to publication, selectivity, and reviewing practices; then after publication as well, with sites like PubPeer offering post-publication review of any journal article by anyone who cares to comment, either with their real names or a pseudonym (Perkel, 2015).

Many commentators online have argued that this is the ideal solution both for the increase in ‘predatory’ journals with poor-to-nonexistent review practices and pay-to-publish models, as well as for the documented failures of traditional peer review to catch errors or misconduct in research. Rather than having just one, two, or at best three selected reviewers, sites like PubPeer open up the review process after publication to the entire community of interested researchers.²⁹ While many publishers have been reticent to allow such commenting on their own sites, in fear of defamation lawsuits, this has given rise to a technological workaround: a simple browser extension allows interested users to have links to PubPeer comments automatically included as they browse through journals online. The liability concerns on the part of traditional publishers are by no means without merit. PubPeer has its detractors as well, who worry that it allows anonymous commenting, and

²⁹ Worse still, in traditional peer review, some journals allow authors to suggest their own potential reviewers: a practice that makes things easier for the editors, to be sure, but has also given rise to troubling ‘cartels’ of researchers who rubber-stamp each other’s papers, and even cases where researchers have given entirely fabricated contact information for their reviewers, and ‘reviewed’ their own work (Barbash, 2015). Even where there is no outright misconduct, the practice of suggesting one’s own reviewers means that they are likely to be either positively disposed to the findings in question, or at least within the same subdiscipline. Picking out obscure problems with a paper, though, often demands a more critical eye, or a different disciplinary perspective, as with the concerns about machine learning in neuroscience discussed above (Combrisson & Jerbi, 2015).

consider it just as likely to become a breeding ground for personal attacks and meritless sniping as for substantive critical review. One of these detractors, Fazlul Sarkar, has indeed gone so far as to sue the site for defamation, arguing that discussions on the site led to the rescinding of a job offer, and demanding its operators reveal the identity of an anonymous commenter who implied misconduct in relation to several of his papers (Oransky, 2014).³⁰

These mediated transformations are taking place across the sciences, but are particularly salient for psychology: from the production of data using Mechanical Turk, to many of the most prominent retractions and cases of misconduct widely publicized online, to the development of new Web-based tools for post-publication review and replication, the cognitive sciences are at the forefront of this shift. In the end, though, all this digital mediation contrasts with the surprisingly sparse literature on cognitive effects of media. One straightforward, normative conclusion I draw is that cognitive science ought to devote more of its efforts toward such research, and toward reflexively applying its findings to the design of its own technologically-mediated practices. This strikes me as all the more important, inasmuch as there are troubling signs that understanding the cognitive effects of technology—indeed, *understanding* cognition in general—may be eclipsed as a practical goal by research oriented toward designing technologies for deliberately creating such effects.

³⁰ PubPeer in fact has a detailed Terms of Service (ToS) intended to pre-empt just such allegations, which includes provisions that comments must adhere to, including refraining from direct accusations of misconduct, or any personal comments about paper authors. Part of the claim in Sarkar's suit is that the specific commenter is liable for defamation, and so their IP address should be revealed; the 'identity' of the commenter was later revealed by Wayne State University through an email exchange as 'Clare Francis,' a pseudonym that has in recent years been bombarding scientific journal editors with claims of image duplication and fraud in their publications. Some of these have led to retractions and findings of misconduct, while other have proven baseless (Yong, Ledford, & Van Noorden, 2013). It is likely that even if PubPeer were forced to surrender this user's IP address, they would have used some proxy service to conceal their true identity. The other substantive claim, however, is that by allowing certain comments to remain on their site, in possible violation of their ToS, PubPeer itself shares in the liability. Prolific, anonymous critics like Clare Francis have thus given rise to a new debate over whistle-blowing, fraud, and the propriety of different venues for scientific debate.

One clear and promising path of research in this regard concerns treatment of psychiatric disorders, particularly through transcranial electromagnetic stimulation, and direct stimulation of the brain with implanted electrodes. Operating under the aegis of the BRAIN initiative, one major program aimed at the development of such technologies is DARPA'S 'SUBNETS' program: "Systems-Based Neurotechnology for Emerging Therapies." As framed by its lead researchers, the program's vision

is distinct from current therapeutic approaches in that it seeks to create an implanted, closed-loop diagnostic and therapeutic system for treating, and possibly even curing, neuropsychological illness. That vision is premised on the understanding that brain function—and dysfunction, in the case of neuropsychological illness—plays out across distributed neural systems, as opposed to being strictly relegated to distinct anatomical regions of the brain. The program also aims to take advantage of neural plasticity, a feature of the brain by which the organ's anatomy and physiology alter over time to support normal brain function. (Sanchez, 2015)

SUBNETS seeks to create a single tool for simultaneously representing and intervening in the plastic brain. It would operate by measuring "pathways involved in complex systems-based brain disorders such as depression, compulsion, debilitating impulse control, and chronic pain" (The Neurocritic, 2013). First it would establish the ability to measure the functioning of brain systems in both normal 'control' populations and individuals with diagnosed neuropsychological disorders, and distinguish between the two through direct measurements, a diagnostic technology which remains elusive despite much talk of the brain in contemporary psychiatry. This would then be translated into "next-generation, closed-loop neural stimulators" that could continuously monitor brain function while attempting to correct pathologies through targeted electromagnetic stimulation (Figure 31, Blaszcak-Boxe, 2014).


The notion of treating the mind by electrically stimulating the brain is by no means a new one. What is new is the proposed technology's distributed nature, and its capacity for simultaneous recording and stimulation. If mental disorders cannot be distinguished by directly measuring one part of the brain or by noninvasive technologies, perhaps a network of many electrodes implanted in various parts of the brain will be more effective. And if so, those same kinds of implanted electrodes may well be able to modify neural firing in addition to measuring it, applying targeted stimulation to many parts of the brain either simultaneously or in patterned ways. This program is of military interest, at least publicly, for its applicability to treatment of the increasing percentage of veterans suffering from mental

disorders, whether from psychological or physiological brain trauma.

Of course it would be of much broader applicability were it to be successful. In principle any mental disorder would be susceptible to

treatment by this technology,

though in practice there would almost certainly be some disorders that were more amenable to modeling and treatment than others. Beyond that, however, if the SUBNETS research program were to achieve even a modest level of success, two major areas of ethical concern would immediately arise. The first would be in the use of such brain stimulation techniques not for treating deficiencies, but for enhancing normal human capacities and extending them beyond typical limits –



In this artist's concept, a miniature electronic device would be placed between a patient's skull and scalp. A series of electrodes placed at varying depths in different regions of the brain would record neurological data. A clinician could wirelessly review the data and communicate with the device to prescribe tailored therapies. /

Figure 31 - DARPA SUBNETS technology rendering

keeping pilots on missions awake for long periods of time, for instance, or allowing for improved aim when firing weapons. Some of these same concerns arise in relation to neurostimulant drugs, which I touch upon in the final section. More uniquely and ominously, however, this technology raises the spectre of real ‘thought control:’ if this could correct complex psychological disturbances such as anxiety disorders or addictions, what if it were targeted at ‘correcting’ normal psychological tendencies which posed problems in institutional settings, such as dissent and defiance? Could this technology be put to use creating ultra-obedient soldiers—or students, for that matter? For now this remains the province of science fiction, but it is a matter of concern when considering any sufficiently advanced technology for the treatment of neuropsychological disorders.

On a more mundane level, gaming is another vast frontier for experimentation with media effects on a mass scale. Apart from using gamers themselves as research subjects, as discussed above, there are games like ‘Neuroracer’ which aim to generate therapeutic cognitive effects, particularly for improving memory in aging populations. Neuroracer is one example that has delivered significant results in major peer-reviewed publications, supporting claims that it can improve ‘cognitive control’ and multitasking ability (Anguera et al., 2013), and is currently under FDA review so that it might be officially deemed a medical device. Other games in development seek to couple with the body more directly, incorporating biofeedback with sensors for heart rate, galvanic skin response, or electroencephalography, often with the intent of training users to control anxiety. There are still other contenders claiming much broader benefits for their ‘brain training’ games, such as Lumosity, but the benefits of such mass-market software are less well documented. If there are any net benefits they may well be less than or equivalent to those of any traditional game (as with the supposed benefits discussed above of ‘first-person shooter’ games). One must also consider the potential opportunity cost of time spent playing these games; no doubt other activities

less readily examined in a laboratory, such as taking a walk or having dinner with friends, may confer cognitive and physiological benefits as well.

Standing at the opposite end of the spectrum from such electronic games claiming to improve brain function and cognitive control, however, are a great many which seem oriented more toward producing compulsive behaviours. The developers of electronic games have always sought to encourage repeat play through design elements such as high scorer rankings, since the days of arcade gaming, when repeat players meant more income for the owners. Electronically controlled gambling machines like slots and video poker similarly promote compulsive play by ‘near misses,’ when in fact all outcomes are fully controlled and no loss is ‘closer’ to a win than any other. Now there are far more advanced ways of monetizing electronic games for the mass market, leveraging mobile payment infrastructure to give very sophisticated games away for free while charging incrementally for virtual items or currency within the game worlds. Such games have a very large userbase with revenue following a power-law distribution: the majority of players try the game and pay nothing, while a minority spends a bit on the game and a tiny percentage spend exponentially larger sums. Designers of such games are often fluent in the neuromarketing buzzwords, searching for the next best way to produce and profit from the ‘dopamine squirts’ of their users.

In practice how these games work is best described in terms of what one executive in this market called ‘fun pain’ (Murphy, 2013): you get users wrapped up in some game activity, like building and operating a farm or a military encampment, but you tie it to some finite virtual resource that regenerates slowly. Eventually you tell players they can’t continue doing the fun thing they were doing, unless they either wait for a while, or pay real money for more of said virtual resource. Another prong of this strategy is that some of the in-game actions required are somewhat repetitive and annoying, but you need to complete them in order to progress in the game—so you can either

grind away for long periods of time, or purchase other virtual items that automate some of the repetition for you. These techniques are developed and refined through A/B testing and comprehensive measurement. Game companies operating in this space are constantly testing different variations on design elements to see what maximizes both the revenue from and the percentage of users willing to spend on ‘microtransactions,’ and releasing new games to sustain the attention of high-investment players who have exhausted the content of the old ones. Absent the possibility of financial reward, these sorts of games are not all that different from gambling machines, in their exploitation of cognitive systems for attention and reward, not to mention in their capacity to produce problematic, compulsive behaviours in a minority of users. When considering their use by children, any electronic games must again be evaluated in terms of opportunity cost and which activities they are displacing. Not without reason, concerns often focus on violent content, but it seems the formal structure of the activity is just as important. Quite irrespective of their content, microtransaction-oriented games on mobile devices are particularly worrisome for their potential costs in terms of both time and money, hence their risks merit special scrutiny by parents.

One may envision the interactions and spending behaviours of game players as constituting a vast potential archive for future sociologists and economists. Their findings could in turn lead to greater refinements still in the abilities of corporations to capture and monetize the attention of the masses. As in many other cases chronicled in this section, media technologies circulate from scientific research sites to mass markets and back again, profoundly influencing the paths of research and the contexts in which findings are disseminated and reinterpreted. There is a strong pre-existing public interest in how technology may be shaping our ways of thinking and living. There are also potent public biases about new media. Such debates are of course longstanding but they take on new vocabularies and valences in an era of digital machines and neurobiological selfhood. On the one

hand, scientific research does seem to support the view that media are reshaping our collective patterns of sensation, perception, and cognition, as McLuhan (among others) contended. The new media which originated in the twentieth century and have only grown more pervasive since, electronic video and interactive computing, may be acting upon our amygdalas to make us subtly more aggressive in our thinking, or altering our capacities for executive function and attention, for better or for worse. Such technologies may support scientific research into cognition; they may also make us unwitting subjects of experiments, or other active attempts to induce cognitive effects on the part of corporate or military interests. The effects are real, and seem poised only to become more so. Their full extent and the normative conclusions to be drawn though remain uncertain, and dependent on the particular medium. How we should govern our conduct and that of our children with respect to technological interaction, and whether we should be fearful or hopeful for a more pervasively mediated future, will continue to be matters for public debate.

Cognitive science can only offer some empirical and theoretical starting points, some guidance as to how things are; we must decide, in light of its evidence among other sources, how things ought to be with respect to technology. Allied with the insights of STS, however, it can help to rule out simplistic unidirectional causal models. It is neither the case that the 'natural' human mind is being uniformly corrupted by the deterministic effects of new media, nor that these technologies are neutral tools which merely aid in the progress of human knowledge while leaving our ways of thinking unchanged. The effects of media are not some new, external danger, nor are they leading inexorably toward transcendence of all human limitation. Instead we are 'natural born cyborgs,' cognizing agents whose mental activity has always been embedded in and distributed across an array of media (Clark, 2003; Hutchins, 1996, 2014). Whether constructing philosophical essays, lines of position, or models of subatomic particles and abstract minds, this external and

tangible mediation does simply record a mental representation, but plays a constitutive role. Digital computing is only the most recent and sophisticated such extension of our cognition. The effects of new media, then, must be evaluated not in terms of how they displace a natural state of being, but to what extent they displace previous mediated forms of attention, reason, and interaction. Beyond that, we must ask whether the collective social effects of these displacements are desirable. This recalls Stiegler's framing of the issue, of 'taking care of youth and the generations,' as cited at the outset of this project. But I have tried to emphasize that it is a far more open question than he supposes as to whether the effects in sum will be positive or negative, and that they must differ as we consider different facets of the recent revolutions in computing and communications.

The other primary line of argumentation in this section and the previous chapters has been that the cognitive sciences are pervasively influenced by digital media, from their origins to their present-day practice. While as we saw the machines of cybernetics were often largely analog or hybrid electromechanical devices, even in those days Wiener and others regarded digital techniques as preferable. Later, the rise of cognitive science, interactive computing, and the ARPANET were closely linked. Now, from the initial stages of production to the dissemination and review of research, the Internet is coming to play an integral role, having travelled full circle from its origins to become a mass medium, and then once again an integral technology for scientific research. It is a venue for publication, simultaneously for formal communications, dynamically revisable consensus statements, and informal 'corridor talk,' but also for experimentation and observation, as a means to access a mass public of research subjects. The roles played by digital media in science more generally are likely to expand further in the future, making it all the more necessary that we seek to better understand how these technologies act upon us. The next section turns from questions of media effects in cognition and science to an analysis of media texts, and to specific claims about the the

interrelation of technology and mental pathology. Could it be that electronic media have such pronounced effects upon attention that they should be held responsible for one of the most prominent psychological disorders in our time?

5

Technology and attentional pathology: the case of ADHD

“It is in the late nineteenth century, within the human sciences and particularly the nascent field of scientific psychology, that the problem of attention becomes a fundamental issue. It was a problem whose centrality was directly related to the emergence of a social, urban, psychic, and industrial field increasingly saturated with sensory input. Inattention, especially within the context of new forms of large-scale industrialized production, began to be treated as a danger and a serious problem, even though it was often the very modernized arrangements of labor that produced inattention.”

— Jonathan Crary (1999, p. 13)

This section constitutes a case study, examining the mediated discourse surrounding one specific form of psychopathology: Attention Deficit Disorder (ADD), and its longtime companion symptom, hyperactivity. Even before its official enshrinement as such within the DSM-III (APA, 1980) and its ascension to become the most common diagnosis of mental illness amongst North American children, the definition of this disorder represented a major locus of public controversy. Named only relatively recently by Canadian psychologist Virginia Douglas (Douglas, 1972), the story of ADD leads back through a much longer history of medical concern with the behavior of children. Foucault links the rise of psychiatry and its conception of mental abnormality to the ‘psychiatrization of childhood’ far more than of adulthood (Foucault, 2006, p. 202). Even in cases of adult behaviour, psychiatric knowledge was typically oriented toward childhood, whether in terms of

causative circumstances or a history of delinquency (Foucault, 1990, 2003). A long view of the pathologization, medicalization, and ‘somatization’ of children’s behaviour forms the background against which I examine contemporary media coverage of attention deficit disorder. Sociologists have held that this disorder is a particularly salient case for the ‘medicalization of deviant behaviour’ (Conrad, 2005), and it has been argued that “diagnostic categories are often fought out as turf battles between medicalizers and their opponents” (P. Brown, 1995, p. 39). The mass media are at once a primary site for this ‘turf battle,’ and actors within it, recruited in a variety of ways. Concern over media effects is manifested in the press alongside a diverse set of other issues: including the reality or constructedness of the diagnosis; the efficacy, side effects, and nonmedical use of psychostimulants; the ethics of ‘drugging children’ and treating their conduct as a matter of medical, pharmaceutical intervention; and the possible etiological roles of other environmental or cultural factors.

Few studies exist of the public portrayal of ADD/ADHD in the media (for one exception see Clarke, 2011), though there is a small corpus of valuable work on neuroscience in the media (Racine et al., 2005; Racine, Waldman, Rosenberg, & Illes, 2010; Robillard & Illes, 2011). The high volume of coverage and the great diversity of concerns expressed within it make attention deficit disorder an excellent case study in the popular communication of scientific knowledge. I argue that while the tone of stories lends some credence to a common perception of news media as controversy-stoking propagators of moral panics, this is not the primary phenomenon of interest. Though relevant, this strikes me as the standard structure of most news reporting, and an unremarkable finding.¹ There may, moreover, be justifiable controversies and panics constructed

¹ Such a tendency can be observed far back in the history of scientific journalism, and has historically been much emphasized in evaluations of its quality. This is typically taken to reflect poorly not only on journalistic practice, but its audiences, with even promoters sometimes expressing a certain disdain for the public. As Dorothy Nelkin recounts the

through the media, and I am personally sympathetic to critiques of the pharmaceuticalization and globalization of psychiatry and ‘mental health’ (Conrad, 2005; Healy, 2002; Kirmayer, 2006; Rose, 2006). Yet my analysis here remains impartial with respect to the claims made by promoters and critics of the psychiatric consensus on ADD. My aim is to highlight the heterogeneity and potential performativity of the popular media discourse on the disorder. I argue that this disorder is necessarily a social, cultural, and technical construct, but not in the same deflationary sense proposed by some anti-psychiatric parties to this public debate (Timimi & Taylor, 2004). As I discuss in more detail below, several such writers have established themselves as ‘passage points’ of a non-obligatory, alternative type, offering claims which circulate through both professional and popular discourse as an interwoven complex of skeptical and debunking arguments: that ADHD is a fiction, largely created by corporations peddling addictive, ‘brain-disabling’ psychostimulants (Timimi & Taylor, 2004, p. 8). Bringing to bear a more sophisticated understanding of social construction, shaped by the many debates on this topic within science and technology studies, I am not interested in debunking the condition as unreal, but instead describing the complex realities at stake in its construction.

I divide this section into two distinct chapters. In the first, after outlining my stance on social construction, I present the historical background of attention deficit disorders. The second then covers the methods and findings of a content analysis conducted along several dimensions of the media coverage concerned with this diagnostic category. My analysis focuses on the overarching concern of this dissertation, media effects, and what I view as a noteworthy aspect of the discourse: it has occasioned an unusual reflexivity, with news outlets regularly voicing concerns that the media

proclamation of chemist Edwin Slosson, first editor of the Scripps Science Service, “it is not the rule but the exception to the rule that attracts public attention. The public that we are trying to reach in the daily press is in the cultural stage when three-headed cows, Siamese twins and bearded ladies draw the crowd to the sideshows” (Nelkin, 1990, p. 44).

themselves are responsible for escalating diagnoses of the disorder. As a typical lede asks: “As U.S. children are exposed to 8 1/2 hours of TV, video games, computers and other media a day – often at once – are they losing the ability to concentrate?” (Elias, 2005). Supported by a relative paucity of solid evidence but a number of ostensibly scientific spokespeople, a number of voices in the press have touted variations on this hypothesis, each having varying degrees of credibility. These articles simultaneously report, amplify, and shape a widespread popular concern that increasing usage and ubiquity of newer media – television, electronic gaming, and the Internet – might be responsible for the substantial uptick in rates of ADD over the past decades (Sclar et al., 2012). Even within the body of articles discussing the disorder in relation to technology, however, these were a minority. After presenting my quantitative findings, I discuss and analyse some of the broader themes revealed through qualitative analysis.

5.1. Historicizing diagnosis

Regarding 'construction.'

The debates regarding social construction within science studies are at this point well-rehearsed, perhaps largely and deservedly set aside. As Ian Hacking puts it, “one person argues that scientific results, even in fundamental physics, are social constructs. An opponent, angered, protests that the results are usually discoveries about our world that hold true independently of society” (Hacking, 2000, p. 4). Social construction should seem an unproblematic thesis when it means we acknowledge the coterie of contextual, cultural human factors deeply influencing scientific labour. Where such analyses can go wrong is in presenting themselves as undermining or debunking the reality of the phenomena which science constructs. In medical diagnosis, it would be quite plausible to discuss the myriad social, economic, and geopolitical factors affecting whether and how one is diagnosed and treated for, say, a fractured toe, or malaria. It becomes less plausible to suggest that somehow these diagnoses are freely invented by the medical establishment, unmoored from any mind-independent state of affairs. Though our understanding of these conditions is shaped by culture, they reflect a reality independent of how we understand it. There are agreed-upon techniques for determining the fact of the matter with respect to these conditions, and a number of different standard methods may typically be brought to bear, if necessary providing a triangulation which resists dissenting opinion. In such diagnostic contexts, social construction remains a productive mode of analysis, but one may analytically separate the ‘internal’ justifying logic of medical diagnosis from the ‘external’ social factors which shape its implementation and interpretation – however intertwined these may be in practice.

When it comes to psychopathology, diagnosis stands on far less stable ground. It becomes substantially more difficult to draw any line between the realms of social construction and empirical objectivity. This is often regarded as a matter of its regime of objectivity being ‘in-the-making,’ unlike in other branches of medicine which appear to the public as already-made, with clear boundaries between their internal, institutional logics and the social world outside. As with most other forms of mental illness, many researchers believe that ADD is a biologically-based disorder with verifiable, distinct markers to be found through brain imaging, but as yet the search for diagnostically useful biomarkers of the condition has borne no fruit (Scassellati, Bonvicini, Faraone, & Gennarelli, 2012). Diagnosis of the disorder remains of the traditional psychiatric type: extensive interviews, typically targeted at children but with questions put to caregivers as well, involving many hundreds of questions about emotions and conduct, covering multiple potential disorders (Shaffer, 2000). These are now often delivered with the assistance of a computer. Symptomatic diagnosis using the DSM-IV may be of the ‘predominantly inattentive’ or ‘predominantly hyperactive’ types of the disorder, and requires the presence of six or more symptoms persisting for at least six months, “to a degree that is maladaptive and inconsistent with the developmental level” (APA, 2000). Symptoms for the inattentive type include difficulties sustaining attention and giving care to details in work, while the hyperactive type involves excessive fidgeting or acting “as if driven by a motor” (*ibid.*).

Such diagnoses based in symptomatology present clear challenges for reliability. Specifying and addressing these is central to the construction of psychiatric objectivity, but by conventional standards many diagnoses fare poorly. In a community sample combining parents and children, the NIMH Diagnostic Interview Schedule has been assessed as having a test-retest diagnostic reliability for ADHD of 0.48 (Shaffer, 2000) – meaning that only about 23% of the variance in results is

shared when this interview is administered twice. Interpreting statistical reliability within psychology is itself challenging, and many researchers insist that this figure is “misleadingly low” (*ibid.*, p.35).² Such diagnostic instruments remain useful, but these measures leave the community of psychiatric expertise open to widespread dissent. When objections are raised that ADHD is overdiagnosed, at present brain images may be deployed to argue in a general way for organic causation, but there is no specific test by which one could demonstrate a brain abnormality in all those diagnosed.

Images like those from X-ray machines, microscopes, and genetic sequencers may come to constitute ‘immutable mobiles’ (Latour, 1988), ‘black boxes’ concealing their process of construction and lending a reality effect to the claims of those actors able to recruit them. As Georges Canguilhem contended, the visibility of germs, even if it requires “the complicated mediation of a microscope, stains, and cultures,” embodies an “ontological representation of sickness” (Canguilhem, 1991, p. 11). The elusiveness of such evidence in the context of most mental illness undercuts any lines we might draw between the hard ‘objective’ core of a condition and the ‘external’ social factors shaping its diagnosis. That division itself is constructed and contested, and I do not wish to suggest that in cases where this has already taken place, the analytic of construction ceases to be useful. Even relatively well-understood biomedical conditions like atherosclerosis have been shown to constitute rich sites of social and ontological negotiation (Mol, 2003). But I argue that in cases of mental phenomena, the public struggle to establish this division shows it to be particularly untenable. There are truly no grounds on which to posit an opposition between their construction within human society and their ‘real’ biological basis; the latter remains obscure and

² To briefly expand on some of the complicating factors here, there is certainly an ‘attenuation effect’ at work, where the very fact of administering a test twice changes the responses given the second time around (Lucas, 1992; Shaffer, 2000). However there is also an enormous difference in test-retest reliability between the DISC parent interview and the youth interview, with the former measured at 0.60 and the latter at 0.10. 0.48 is the combined measure for ADHD in a community sample, which increases to 0.62 in a clinical sample, and reliability can be further improved by incorporating measures of generalized functioning such as the Children’s Global Assessment Scale (Schwab-Stone et al., 1996).

itself deeply social. Even if biomarkers for the disorder are eventually found through brain imaging, causal interpretations remain underdetermined, given that connections within the brain are shaped through our interactions with the world. Widely accepted only since the 1990s, scientific conceptions of neuroplasticity are becoming incredibly influential, especially in debates over the effects of digital media (Choudhury & McKinney, 2013; see also Rose & Abi-Rached, 2013). Yet this research raises more questions than it answers. Are patterns of media usage and dysfunctions in social settings *caused by*, or *causes of* brain abnormalities?

The classification of individuals having ADD/ADHD is not a matter of ‘natural kinds,’ but of *human, interactive* kinds (Hacking, 1995). Psychology is performative discourse *par excellence*, in its “reiterative power...to produce the phenomena that it regulates and constrains” (Butler, 1993, p. 2). It is not the case that ADHD has always existed as such, simply waiting to be discovered like *Mycobacterium Tuberculosis*; conversely, children were not free, happy, and normal until the arrival of imperialist psychiatry and its invented diagnosis. (Both positions, however, find voice today within the popular press.) Diagnosis is not typically imposed by any one privileged actor, though again this implicit sociology shapes popular critiques of medicalization. Rather, the construction of ADD, as with Hacking’s examples of multiple personality disorder and autism, involves a complex multi-directional interaction between cultural problematization of behaviour, communities of parents, educators, and physicians, the pharmaceutical industry, and a mediated public discourse. The fact of being diagnosed fundamentally reshapes one’s self-conception and social sphere, transforming an ‘unorganized’ and heterogeneous collection of worrisome experiences into the symptoms of an ‘organized illness’ (Balint, 2000 [1957]). Whatever one’s view of the disorder’s biological basis, or lack thereof, the end result of classification is ‘making up people’ (Hacking, 1986): articulating a

subject-position which did not previously exist, that of the patient with a medically treatable attention deficit.

The articulation and individuation of this ADD-subject involves simultaneous definition of the normal and the pathological. The normal is not a 'static and peaceful, but a dynamic and polemical concept' (Canguilhem, 1991, p.146), not an enduring object in itself but the "effect obtained by the execution of the normative project" (*ibid.*, p.149). Nowhere is this clearer than in the context of ADD. To define certain patterns of action and perception as abnormal and deserving of medical intervention is, by the same gesture, to outline its obverse realm of normal human variation. The 'normal child' is not an entity which preexists, but emerges from, the discourse and techniques of psychiatry. The studies of media coverage which follow are an attempt to glimpse this normative project in process.

Departing from this conception of ADD/ADHD, both mediation and social construction are of dual significance. The media are a forum for presenting and contesting differing conceptions of the disorder, and thus a possible vector for the recruitment of scientific allies and for looping effects on self-conception. But moreover, the media are a potential culprit, often flagged anecdotally and within the press as wholly or partially to blame for the rise in diagnosis. Hence the popular media are at once centrally involved in the 'execution of the normative project' with respect to ADD, and a common object of concern for a kind of folk epidemiology of the disorder.³ Likewise, the media are not only integral actors within this project of social construction, but sites for popular discussion and interpretation of the very ideas of constructivism. The escalating diagnoses of

³ What follows can indeed be seen as a study in the folk or popular epidemiology of attention deficit disorder, as defined by Phil Brown in his work on toxic waste exposure: an inquiry into the "process by which laypersons gather statistics and other information and also direct and marshal the knowledge and resources of experts in order to understand the epidemiology of disease" (P. Brown, 1987, p. 78).

attention deficit disorder occasion extensive discussion of the question which I follow Hacking and Latour in regarding as poorly formed: is this a real disorder which we are simply getting better at discovering, or are all of these new cases being invented by a corrupt psychiatric-pharmaceutical apparatus?⁴ In this instance as in others we can begin to move beyond this dilemma by gaining a longer historical view on the diagnosis.

⁴ This is also widely debated in relation to autism (Eyal, 2010). A crucial difference with the case of autism is the existence of pharmaceutical treatments whose marked effects on ADHD have long been observed and documented: the psychostimulants, whose history I discuss below. The wide range of patented and generic drugs currently approved for the treatment of attention deficit disorder implies that the pharmaceutical industry has both incentive and credibility to take an active role in the shaping of public understanding. This is less true in the case of autism, for which there are few drug treatments approved, and minimal evidence of their efficacy in meta-reviews of research (Broadstock, Doughty, & Eggleston, 2007).

From 'moral treatment' to 'moral control.'

I focus principally here on the ways in which our concepts of illness and abnormality are constructed in the public sphere, and less on the construction of *medical knowledge* within the expert community (Brown, 1995, p.37). Evidently the two processes are connected, however, particularly in cases like this where the status of the medical knowledge remains a contested matter of public concern. Before proceeding to the results of my media research, I offer some context from the prior literature on the social history of this medical knowledge itself. This is especially valuable here, because the kinds of themes and concerns expressed long ago in the 'prehistory' of attention deficit disorder recur perennially within contemporary media coverage. I hesitate to call this a history *of* the disorder, inasmuch as my stress will be on its discontinuity and malleability over time.

Until quite recently, the emphasis in diagnosis was firmly on children and on hyperactivity. Difficulties in sustaining attention were a prominent feature of the cases described in the early literature, but they were not typically taken to be the most salient aspect of the disorder. The very earliest work to discuss symptoms comparable to ADHD was an exception, however – a chapter focused specifically on 'morbid' abnormalities of intention in the Scottish physician Alexander Crichton's volume, *An Inquiry into the Nature and Origins of Mental Derangement* (Crichton, 1798). Crichton's discussion of a condition which appeared in childhood as an incapacity "of attending with constancy to any one object of education" (*ibid.*, p.271) has been cited by historians as the first recorded mention of ADHD-like symptoms, and in Whiggish accounts intended for medical audiences is even labeled as a depiction of the "same disorder as defined in the current DSM-IV-TR criteria of ADHD" (Lange, Reichl, Lange, Tucha, & Tucha, 2010).

Another early point of reference is the nineteenth-century German physician Heinrich Hoffmann, in whom we find a striking juxtaposition of medicalization and moralization. Hoffmann was a successful psychiatrist at the mental hospital in Frankfurt, supposedly known for his efforts in improving the treatment of patients there (Lange et al, 2010, 243); he was also the author of a famous series of illustrated ‘cautionary tales’ for children, *Der Struwwelpeter* (“Slovenly Peter”), in which the bad behaviour of children led to all manner of catastrophic results, including occasionally their deaths. Amongst these was the story of ‘Zappelphilipp,’ or ‘Fidgety Philip,’ who “wiggles and giggles” and won’t sit still at dinner, provoking his parents’ anger and eventually destroying the entire table setting in a fit. Again, particularly in progressivist historical accounts within the professional medical literature, this poem has been cited as an early symptomatology of ADHD and, indeed, as solid evidence that “the diagnosis of ADHD is not an ‘invention’ of modern times” (Thome & Jacobs, 2004). The character has even become a mascot of sorts on Facebook for an awareness program of Children and Adults With Attention Deficit Disorder (CHADD), the largest advocacy group for the disorder in the United States (National Resource Center on AD/HD, 2013). Hoffmann’s professional life makes this historical episode all the more intriguing, yet if anything the poem illustrates that he diagnosed ‘Philip’s’ behavioural symptoms as those of a “naughty, restless child... rude and wild,” *not* as indicators for medical-psychiatric intervention.

The more direct scientific starting point for histories of attention deficit disorder is a series of lectures given before the Royal College of Physicians in 1902, by George F. Still. Crichton’s discussion of attention was part of a philosophical anthropology, a systematic discourse of mental faculties and their ‘derangements,’ while Hoffmann’s was a fable, the very paradigm of *moralizing* discourse. With Still’s studies we confront an approach that, however unfamiliar, is profoundly contemporary – concerned no longer with the abstract and general systematization of the faculties,

but as Foucault described modern psychiatry, with establishing authority over delinquency, deviant sexuality, and dangerous criminality by establishing their “antecedents below the threshold” (Foucault, 2003, p. 19) and their physiological, somatic underpinnings. The medicalization of ADHD is currently thought to correspond with a shift away from ascriptions of moral responsibility: diagnosis changes an individual’s status from ‘bad’ to ‘sick.’ Or, as the critics of medicalization argue, it provides the spurious excuse of illness, and makes individuals targets of medical control. In Still’s lectures, however, we find the first of many uneasy hybrids, blending somatization – in an emphasis on brain dysfunction and clinical intervention – with the more established conceptual schema of moral responsibility and volitional control.

He proclaims this a matter of “very real -- shall I say social -- importance” (Still, 2006, p. 126), detailing a number of cases of children who exhibit symptoms of impulsivity, ‘passionateness,’ ‘lawlessness,’ violent outbursts, and inappropriate self-gratification. He does highlight their “quite abnormal incapacity for sustained attention” (Still, 2006, p. 133), but only in his third lecture. The major contention of his lectures is that these children constitute a distinct group from those with intellectual impairments, and should not be mistaken for such.⁵ He clearly describes something akin to the symptoms of attention deficit-hyperactivity disorder. It is, however, impossible to view Still, nor any other figure within this history, as the ‘maverick discoverer’ of some mind-independent reality (A. Lakoff, 2000, p. 154; cf. Lange et al., 2010). What reality they discover is within the mind, first of all, but Still’s framing of his material in terms of “defective moral control as a morbid

⁵ As he argues it: “defective moral control as a morbid condition in children [is] a subject which I cannot but think calls urgently for scientific investigation. It has long been recognized that such a deficiency may occur in association with those disorders of intellect which are ordinarily recognised as idiocy, imbecility, or insanity, and I suppose no one doubts the morbid nature of the moral defect in these cases, whether it be regarded as dependent upon the intellectual failure or not. But there are other cases which cannot be included in this category — children who show a temporary or permanent defect of moral control such as to raise the question whether it may not be the manifestation of a morbid mental state, but who nevertheless pass for children of normal intellect; it is this condition in particular which seems to me to call for careful observation and inquiry” (Still, 2006, p.126).

condition” is, despite certain familiar elements, the product of a profoundly different world. It is borne of an irreducibly different system of classification. This should not be a controversial claim: a different social context leads clinicians not only to construct their diagnoses in distinct ways, but leads them to focus on different sets of behaviours as potentially diagnosable.

Adam Rafalovich plausibly links Still’s problematic less with the diagnosis of ADHD which was to follow from it, and more with those of ‘idiocy and imbecility’ which preceded it (Rafalovich, 2001). Idiocy already had a long history of common usage to denote diminished mental capacity, but since the nineteenth century, medicine had been defining it as a condition of extreme deficiency involving substantially reduced capacity to control one’s actions and care for oneself. Imbecility, then, was articulated as a less pronounced form of mental disability, “in a partial effort to provide clarity to the diagnosis of idiocy” (Rafalovich, 2001, p.99). George Still sought to promote a special diagnosis for his cases on the grounds that they constituted a distinct category of imbecile, “too intelligent to be considered ‘idiots’ ... and too young to be viewed as ‘criminal minds’” (Mayes & Rafalovich, 2007, p. 438). Their imbecility was seen as a specific deficiency or delay in the development of their moral, as opposed to intellectual, capacities; their inability to restrain their impulses toward self-gratification, the cornerstone of maturation into adulthood, made them deviations from normal development and subject to confinement as potential risks to society. In this sense, Still’s understanding of these children was of a piece with the ideas of ‘moral insanity’ formulated by J.C. Prichard (Augstein, 1996), and further developed by Emil Kraepelin (Kraepelin & Diefendorf, 1915), to be differentiated more by their youth than by any categorical difference. These

were a new medical discovery, Still claimed, and though he was admittedly unclear on the precise brain abnormalities involved, he was confident that these could be found.⁶

The notion of organic causation is further entrenched subsequent to the observation of comparable behavioural abnormalities in children suffering from brain damage after an outbreak of *encephalitis lethargica* (Ebaugh, 1923). In the work of British psychiatrist Alfred Tredgold (1922), the same basic set of concerns from Still's lectures are linked together with these observations and further quasi-phrenological investigations, outlining the somatic parameters by which this disorder would be known for much of its history. Tredgold helps to establish a consensus that Still's 'defects of moral control' were grounded in abnormalities of the brain, but does so in a curious fashion:

After the passing of the Education Act of 1876, making attendance at public elementary or other schools compulsory, it gradually became apparent that a group of children existed who were so far mentally defective that they could not be satisfactorily taught in the ordinary public schools, but who were not sufficiently defective to be certified as imbeciles or idiots under the Idiots Act of 1886. (Tredgold, 1922, p.174)

Tredgold, as well, is concerned with the question of 'imbecility' with respect to the regulation of behaviour, distinct from the forms of intellectual impairment recognized by the law as 'idiocy.'

Without apparent dissonance, this disorder is recognized in this early phase as a disorder of adaptation, a problem made manifest not by any dramatic *change* in the brains of the nation's children, but by a shifting institutional context, in which the brains of children newly *become* a matter of concern following the establishment of compulsory public education. Although relatively few survivors of encephalitis actually fit modern diagnostic criteria for ADHD, the notion of

⁶ To this end, he was particularly interested in the presence of abnormal skull shapes and other phrenological 'stigmata of degeneration,' taking family histories to investigate the heritability of this moral degeneracy (Still, 2000, p.130). And though these implications are not drawn out extensively in his lectures, the invocation of idiocy, imbecility, and degeneration necessarily ties his claims to the evolutionary discourse of eugenics and the criminology of Cesare Lombroso (Comstock, 2011, p. 48). Foucault also describes the early discourse of idiocy and imbecility in detail, recounting numerous cases from Édouard Seguin and other nineteenth-century French psychiatrists featuring symptoms of inattentiveness and impulsivity: children "uncontrollably lively; climbing like a cat, slipping away like a mouse" (2006, p.217). This further problematizes the drawing of any firm historical line of 'discovery' for ADHD, suggesting rather a story of the gradual nosological disaggregation of intellectual from 'moral' idiocy.

'postencephalitic behavior disorder' worked to "buttress the notion that there was a link between both severe brain damage and severe behavioural disturbances and, by extension, mild brain damage and mild behavioural disturbances" (Mayes and Rafalovich, 2007, p.441). This link was likewise supported with reference to skull measurements of non-encephalitic children presenting similar symptoms. Through analogical extension, then, the cases of children with recognized brain injury and co-occurring symptoms of impulsivity and inattentiveness were taken to justify the notion that some brain damage was likewise present in those with only the symptoms and no apparent organic lesions. This view was eventually codified with Alfred Strauss' term 'minimal brain damage' (Strauss & Lehtinen, 1947), which remained one of the most enduring diagnostic categories until the invention of ADD/ADHD, coexisting with those focused on hyperactivity and eventually metamorphosing into 'minimal brain dysfunction' as the presence of physical lesions became more doubtful (Kanner, 1972; A. Lakoff, 2000, p. 152).

Hyperkinesia, psychoanalysis, and psychopharmacology.

What are the causal relationships between these children's brains, their ostensibly deviant behaviours, their upbringing, and the broader culture? Some place the emphasis squarely on the brain while absolving culture and parenting – this has been highlighted by Andrew Lakoff, among others, as a major factor in the expansion of ADHD diagnosis (2000, p.166). But even within the medical community, medicalization is contested, and this neuro-reductive narrative has had shifting fortunes. In the first major publications on the topic, such as those of Still and Tredgold, we see a straightforward privileging of biology over environment, presented through a collection of case studies of children typically conducted in custodial educational institutions. These routinely adopt

surprising descriptions of this biological disorder in quite traditional terms of ‘morality’ and ‘character,’ while invoking deprecated medical terms with eugenic resonances (idiocy, imbecility, *etc.*). Theirs is “an illness that is not an illness since it is a moral fault” (Foucault, 2003, p.20) and a sign of embryonic criminality, yet its basis is seen as exclusively biological.

Subsequently, from the 1940s through the 1970s, we can observe the rise and fall of a more generally psychoanalytic and dynamic conception, where although brain function is recognized as correlated with mental illness, social, environmental, and parenting factors are privileged. With Charles Bradley’s work and the rise of pediatric psychopharmacology, discussed below, we can see signs of this uneven transition in progress. As Lakoff argues, “the resurgence of descriptive nosology in child psychiatry, which was coemergent with the rise of ADD, can be traced in professional journal articles” (2000, p.154). We observe a shift from the era of analytically-influenced case studies with extensive talk of “overprotective mothers, those reticent weaners responsible for a variety of cold war pathologies” to the contemporary approach, focused on statistical measuring of populations, clinical trials of pharmaceuticals, and the establishment of standardized diagnostic questionnaires (Lakoff, 2000, p.155). This involves a concurrent shift from inpatient to outpatient treatment, and a substantial expansion in funding by government bodies and the pharmaceutical industry. The development of the ADHD diagnosis was thereby manifestation and signal of this wider transition in North American pediatric psychiatry, not only contemporaneous with it but also its most controversial public face. The 1980 publication of the American Psychiatric Association’s third revision of its Diagnostic and Statistical Manual of Mental Disorders, in which ADHD was first labeled as such, marks the culmination of this transition within psychiatry, and opens up a wider debate *about* its new system of diagnosis.

The disorder-in-becoming received limited attention throughout the interwar years, with some authors carrying its implications further into the territory of eugenics. Tredgold insisted on the strong heritability of the condition, but at the same time that it was not restricted to a particular social class, and neither he nor Still alluded to questions of race. Others gave it a more fully ‘Social Darwinist’ treatment, linking the over- and under-development of certain brain areas to “certain plus and minus members of the species” (Kahn and Cohen, 1934, p.750; *cf.* Mayes & Rafalovich, 2007, p.442) possessing greater or lesser attentional capacities. While raising these forms of deviant behaviour as potential domains of medical knowledge, they still offered no distinct course of treatment. The situation remained largely as Foucault described it in the nineteenth century: with their deficiencies seen as matters of incomplete development, the course of treatment was to be “no different in kind to that given to any child,” to “impose education on them.” Therapy in this era was thus “pedagogy itself,” a ‘special education’ which constituted a deeper, more structured and searching modality thereof, “but a pedagogy all the same” (Foucault, 2006, pp.209-210). A sea change began with the work of Charles Bradley, Margaret Bowen, and Maurice Laufer at the Bradley Home in Rhode Island, an institution for children with particularly challenging ‘emotional problems.’ Influenced by the work of Kahn and Cohen, numerous children with similar complexes of deviant behaviours were observed at the Bradley Home, and it was first named as a disorder there by Laufer: “hyperkinetic impulse disorder.”⁷

It was also at the Home that Charles Bradley, nephew of its founders, began the first experimentation with the use of stimulant drugs in these children – several years in advance of the more widely discussed point of departure for psychopharmacology, namely the discovery of

⁷ This preceded the nomenclature of ‘minimal brain damage’ slightly (Mayes and Rafalovich, 2007, p.442). For the most detailed historical discussion of the overall treatment milieu of the Bradley Home, see (Bromley, 2006).

chlorpromazine (López-Muñoz et al., 2005). As with Rhône-Poulenc's search for novel antihistamines, this landmark moment in psychiatry occurred largely by chance, in this case as an odd supplement to psychosurgery. Children in the Bradley home were often subject to a painful procedure known as a pneumoencephalogram – involving the draining of cerebrospinal fluid from around the brain to facilitate X-ray imaging – after which they complained of headaches. Bradley then began giving them Benzedrine (*dl*-amphetamine) “which he hoped would stimulate the choroid plexus – located in the ventricular system of the brain – to produce spinal fluid and thereby reduce pressure on the children's sinuses,” relieving their headaches (Mayes and Rafalovich, 2007, 442). He quickly found that their ‘hyperkinetic’ symptoms were much improved (and thereafter seems to have set aside the question of headaches). He set up an informal single-blind study by also giving Benzedrine to a control group of children who had not recently undergone the procedure. Finding they too were improved, he and Bowen conducted a more formal three-year clinical study of the drug in a sample of one hundred children with various forms of “maladjustment or delinquency” (C. Bradley & Bowen, 1941, p. 94). The study was not generally controlled with placebos, instead they only occasionally tried discontinuing the treatment or substituting for placebo briefly (with reports of ‘relapse’), while arguing that ‘specially prepared placebos’ were typically unnecessary since the children were invariably taking some form of medication daily which could serve as a placebo treatment (*ibid.*).⁸ I quote from their study at length below and reproduce some of the study's tables of results (Figure 32), not only to indicate the findings in detail, but to illustrate the style of his data

⁸ As David Healy notes, in spite of the common perception amongst psychiatrists that randomized placebo-controlled trials had been successfully demonstrating the efficacy of psychotropic drugs since the 1960s, these were relatively rare until they were instituted in the mid-1980s by the intervention of the Food and Drug Administration in the United States (2002, p.307). While this made it a challenge to know whether a drug really worked or not, Healy and others also emphasize that over-reliance on randomized clinical trials with placebo can mislead (Shorter, 2008).

collection and taxonomy. The authors were by no means conducting a study of ADHD in the contemporary sense, nor did they adopt a straightforward conception of biological causation.

TABLE III
Types of Behavior Response Noted in 100 Maladjusted Children Receiving Amphetamine Sulfate

Behavior subdued.....	54
No change in overt behavior.....	21
Generally stimulated.....	19
School performance accelerated without other behavior change.....	6
Total.....	100

TABLE V
Therapeutic Effect of Amphetamine Sulfate in 100 Maladjusted Children Grouped According to Clinical Diagnosis

Clinical Diagnosis	Therapeutic Result			Totals
	Im- proved	Un- affected	Worse	
Behavior disorders of psychogenic origin.....	43	15	3	61
Behavior problems in children with convulsive disorders.....	12	0	3	15
Schizoid personality.....	5	2	0	7
Behavior disorders dependent on structural neurological defect.....	3	3	0	6
Maladjustment due to intellectual deficiency.....	3	1	1	5
Specific reading disability.....	4	0	0	4
Post-encephalitic behavior disorder.....	2	0	0	2
Totals.....	72	21	7	100

Figure 32 - Benzedrine study results

Instead, the text is marked throughout with a concern for sociocultural factors influencing the disorder's 'psychogenic' course of development. The relatively simple collection of population data was supplemented with extensive qualitative descriptions of case studies:

When amphetamine sulfate in daily morning doses of 20 mg. orally was started there was immediately a definite change in his behavior. On the ward John was much quieter, and none of his usual hyperactivity was noted. He was prompt for meals and school and became pleasant and congenial with children and adults. He cooperated well to all matters of routine; no longer restless but applied himself to daily tasks. In the classroom he accomplished a great deal each day and showed excellent initiative . . . His mother, who did not know that John was receiving medication, commented during a visiting period that she had noted great improvement and wondered to what it might be attributed. Following discharge from the hospital John received amphetamine sulfate periodically. At times it produced dramatic improvement in his behavior in school and at home. At other times when unfavorable environmental situations existed little effect was observed. (Bradley & Bowen, 1941, p.96)

This was presented as typical of the ‘behavior subdued’ category. Other boys’ reports only indicate their being ‘generally stimulated,’ but in a manner which improved overall social functioning and any ‘schizoid’ or depressive tendencies. Both responses were seen as bringing the children closer to normal social functioning, corresponding with distinct hyperactive and withdrawn expressions of the same underlying maladjustment.

Bradley and Bowen noted that few of the children experienced an improvement in school progress without other change in behaviour, indicating again that the change was a matter of broader conduct rather than exclusively in learning or attention. Twenty-one of them seemed resistant to the effects of the drug, apparently exhibiting neither therapeutic nor even ordinary stimulant effects from doses of up to 30mg of amphetamine, while seven of the hundred exhibited ‘retrogression’ and ‘accentuation’ of their “excessive activity or irritability” in response to the drug. The authors of the study seemed somewhat surprised by the paradoxical ‘subduing’ effect of stimulant drugs on the children, and noted that continued drug therapy had benefits but was often disrupted once they had returned to their communities: “in most private homes and average communities, situations frequently arise which are sufficiently disturbing emotionally to offset temporarily the beneficial effect of the drug” (Bradley & Bowen, 1941, p.442). Though not wholly conclusive by any modern standards, the Bradley Home research was taken as demonstrating remarkable levels of efficacy for psychostimulants in treating these disruptive children, many of whom had proven resistant to behavioral therapies. The challenge, then, was to construct some plausible account of *why* they were so effective.

Often overlooked by internalist accounts emphasizing progress toward neurobiological psychiatry (Lange et. al, 2010), critical historians have remarked upon the shifting mixtures of biological and environmental causation found in these early studies of attentional abnormalities

(Lakoff, 2000; Rafalovich, 2001; Mayes & Rafalovich, 2007; Bromley, 2006). The work of Bradley, Bowen, Laufer, and their collaborators represents a turning point in this regard, predating and prefiguring the opposition between neurobiological and psychoanalytic approaches. Their publications present some of the first detailed claims for somatic mechanisms underpinning ‘hyperkinetic impulse disorder,’ suggesting that it involved dysfunctions in a portion of the brain known as the diencephalon, which amphetamine could treat by somehow preventing “the cortex from being flooded by streams of unmodulated impulses” (Laufer & Denhoff, 1957, p.46). At the same time, however, they retain the basic framing of Adolph Meyer’s analytic psychiatry, insisting on the ‘psychogenic origin’ of these behaviours in maladjustment, ‘ego disturbance,’ and unhealthy parenting – be it the overprotective mother or the permissive father. On their view, amphetamines could be a helpful adjunct to psychotherapy, subduing undesired behaviours and facilitating the therapeutic interaction, but offering no substitute for its “fostering of a basic inner change” (Laufer & Denhoff, 1957, p.44; *cf.* Mayes & Rafalovich, 2007, p.445). Such juxtapositions indicate that ideas of neurobiology and pharmaceutical treatment were by no means mutually exclusive with those of psychoanalysis and sociocultural etiology. Demonstrating the efficacy of drug treatment was a crucial factor in the reconceptualization of this ‘defect of moral control’ as fundamentally caused by a brain disorder, but by no means did it render this transition inevitable.⁹

⁹ On the contrary, it has been argued that this transition in its sharpest form was limited to Anglo-American psychiatry, and that in contexts like France (Vallée, 2002) and Argentina (A. Lakoff, 2005), the discourse remains more heterogeneous in its blending of the psychoanalytic and psychopharmaceutical approaches. In those nations, psychiatry seems to retain much the same attitude described above, viewing drugs as supplementary facilitators of the primary analytic treatment. There has been considerable resistance to the APA’s manuals since the DSM-III which introduced ADHD, with child psychiatrists in France instead employing the CTFMEA, a manual intended to address “structural psychopathological considerations” rather than “isolated symptoms” (Mises et al., 2002). This has in turn provoked some interest in the Anglo-American media, with one viral article purporting this manual accounts for why “French Kids Don’t Have ADHD” (Wedge, 2012).

Regimes of treatment, regimes of diagnosis.

Efficacy, after all, is a complex question. From the 1960s through to the 1980s, there was a proliferation of clinical research regarding the value of amphetamines as well as the newer Ritalin (methylphenidate) in treating hyperactivity, coextensive not only with the broader rise of psychopharmacology (Healy, 2002) but with an ‘epidemic’ peaking around 1971, of widespread adult amphetamine use driven by liberal prescribing as ‘diet pills’ or as ‘tonics’ for myriad minor complaints.¹⁰ A consensus was thus established regarding both the effectiveness and the risks of psychostimulants in medicine, even before the origin of ADHD proper, all within the framework of a waning analytically-oriented psychiatry employing the category of hyperkinesis. Talk of ‘attention deficits’ had been present in the literature for some time, but it had not yet become the hallmark of the disorder. Before addressing the DSM-III and the coining of ADHD, I wish to draw out two themes from this period, themes which remain of particular relevance to the framing of issues within contemporary discourse.

First, the establishment of consensus around the efficacy of psychopharmaceuticals involved a realignment of the normative project regarding the nature and primary object of the drug effects. That they had some effect was incontestable, but its mechanism and meaning remained uncertain. Bradley and other earlier researchers insisted that the stimulants operated at a behavioural level, producing a “sense of stimulation, well-being, and confidence” such that psychic conflicts, “though still present, are no longer irritating and distressing” to the extent that ‘abnormal or unacceptable’

¹⁰ These patterns of prescribing gave rise to a ‘speed culture’ of nonmedical use of amphetamines diverted to the black market, and then to a backlash, moral panic, and subsequent strengthening of regulation: for detailed examinations of this history, see ((Rasmussen, 2008); (Moon, 2009)). In the course of this backlash, while the widespread usage of amphetamines as ‘diet’ or ‘pep’ pills was gradually suppressed by the conjoined efforts of the U.S. government and professional medical associations, even sharp critics of the drugs often condoned their prescription to ‘hyperkinetic’ children (Moon, 2009, p. 279).

conduct results (Bradley & Bowen, 1941, p.101). By the 1970s, however, researchers were insisting that drugs' effects were *not* to stimulate or calm the child, but rather to “stabilize the brain” and normalize its functioning (Comstock, 2011, p.56). Those references to matters of social context which were so pervasive in the earlier literature were gradually expunged, and references to external norms of behaviour were supplanted by references to internal norms of neural function and self-government.¹¹ Correspondingly, responses to the drugs began to function as diagnostic criteria in and of themselves through an act of clinical “reverse engineering” (Rafalovich, 2004, p.61): those with normal brain function were said experience the stimulants' ordinary effects of increased activity and euphoria, while those who exhibited paradoxical ‘subdued’ responses were defined as pathological. Drugs were then said to augment ‘executive function’ in those for whom it was deficient, a condition which encompassed a spectrum of abnormal behaviour far greater than before. The population in question was thus no longer limited to those children who posed an individual danger of criminal deviance due to a deficient capacity for ‘moral control;’ the diagnosis expanded to encompass a broader realm of those with an inability to ‘stop, look, and listen,’ with or without hyperactivity (Douglas, 1972), and eventually in adulthood as well. The spectre of criminal deviance nevertheless haunts contemporary discussions in the form of a statistical propensity amongst those with ADD, said to be reduced by psychotherapy and pharmaceuticals, but exacerbated by potential underdiagnosis and undertreatment.

The second important development in this era was of an organized opposition to the nascent psychopharmacological consensus, not only with respect to ADHD, but far more broadly.¹²

¹¹ This coincides with a wider shift toward a conception of the biogenic amine system as the “obligatory passage point of all accounts of mental disorder,” typified by the dopamine and monoamine theories of schizophrenia and depression (Rose & Abi-Rached, 2013, p. 37).

¹² For more on the history and sociology of the anti-psychiatry movement, see (Dain, 1989; Crossley, 1998). Its major figures were Thomas Szasz, coming from the American political Right, and R.D. Laing, from the British political Left

Following the expansion of government funding for studies of psychostimulant efficacy, before ADHD was even formalized as a diagnostic category, there was substantial public controversy about the medicalization and pharmaceutical treatment of children's behaviour. Early media stories expressed concern that parents were being coerced into giving their children behaviour-controlling drugs (Maynard, 1980; Hentoff, 1972), provoking a 1971 Congressional hearing into the matter in the United States (Mayes and Rafalovich, 2007, p.450), but no legislative action. George Lucas' 1971 film *THX-1138* was another signal of developing alarm regarding drugs as mechanisms of discipline, depicting an underground city in which sexual intercourse is illegal and compulsory pills suppress citizens' emotion and desire. Controversy peaked in 1975, when in the same year as Milos Forman's anti-asylum fable *One Flew Over the Cuckoo's Nest* won five Academy Awards, journalists Peter Schrag and Diane Divoky published their *Myth of the Hyperactive Child and Other Means of Social Control* (Schrag & Divoky, 1976), sociologist Peter Conrad offered his critical analysis of 'The Discovery of Hyperkinesis' (Conrad, 1975), and Benjamin Feingold argued that hyperactivity was real, but in fact caused by food additives (Feingold, 1975), a claim which persists in popular discourse (Smith, 2011).

Across all of these works, one finds outlined the basic tenets of a critical counter-discourse which endures today. It contends that ADD/ADHD is the product of an unchecked expansion of medical power, peddling dangerous drugs as a form of chemical behaviour control which (at best) can only mask the disorder's true causes. Not only the problematic behaviours but their 'subdued' responses to stimulant medication constitute durable layers of reality within the construction of the disorder. Yet the ways in which these are to be assembled together remain profoundly contested and underdetermined by the available evidence. The anti-psychiatry movement cooled by the 1980s, but

(though the former would come to disavow the label, which he viewed as emanating from the latter: (Szasz, 2009)); Foucault's work on psychiatric power emerges from this same period, and manifests clear affinities (Sedgwick, 1981).

its strategies of dissent persist in contemporary ADHD debates. Likewise, psychiatry has already formulated its tactic of resistance to claims of ‘drugging children,’ as discussed above: the technique of drug therapy is said not to be oriented toward producing compliance with some external norm of conduct, but toward bringing about stabilized functioning in relation to an internal, neurobiological norm. As we shall see, this shift within the professional discourse remains a matter of fierce contention in public debates, typified by two central metaphors: are psychostimulants ‘correcting an imbalance’ within the brain, or are they a form of ‘chemical straitjacket’ for behavioural control?

This regime of treatment and its critical counter-discourse are thus not products, but precursors of the ADHD diagnosis. The origins of the disorder as such are to be found in the movement of the American Psychiatric Association away from psychoanalysis, and toward a neo-Kraepelinian model of descriptive nosology: diagnosis based neither on ‘psychogenic’ or biological causal factors, but on purely descriptive categories, clusters, and checklists of observable symptoms (A. Lakoff, 2000). With the appointment of Robert Spitzer to the steering committee for the third revision of its Diagnostic and Statistical Manual of Mental Disorders, the APA seemingly marks the end of an era of uneasy mixtures. Spitzer was best known at the time for leading the efforts to remove homosexuality as a disorder from the DSM-II classification, a landmark campaign of *demedicalization* which succeeded in 1973. It seems there was little competition for the chair of the revising committee, and upon his appointment Spitzer created a task force which shared his disdain for psychoanalytic etiologies. Drawn mainly from Washington University in St. Louis, the group instead favoured standardized lists of symptoms with minimum quantitative thresholds for diagnosis, as proposed by John Feighner in 1972 (Kendler, Muñoz, & Murphy, 2009). Historians have offered valuable discussions of the ‘horse-trading’ involved in Spitzer’s revisions, which provoked substantial resistance from the large population of psychoanalysts still within the APA,

triggering compromises exemplified by the definitions of ‘major depressive disorder’ and ‘dysthymia (or depressive neurosis)’ (Shorter, 2008, pp. 153-163; *cf.* Healy, 2002). The case of ADHD, however, seemed to pose fewer obstacles to Spitzer’s vision of replacing analytical hermeneutics with measurement by reliable checklists.

Partly in deference to the continuing clout of the psychoanalysts, and partly due to the paucity of evidence, this method of measurement was in principle to be agnostic with regard to the real origins of mental illness. Its classifications were an attempt to lend greater scientific objectivity to diagnosis of disease entities derived from a widely negotiated consensus of clinicians, not on the basis of but in the absence of consensus regarding their causes. Still, the working groups appointed by Spitzer included many proponents of ‘biogenic’ models of mental illness, and the Child Disorders Committee was no exception, featuring a number of researchers who had written on hyperactivity (a topic largely ignored by psychodynamic theorists), including Paul Wender, Dennis Cantwell, and Stella Chess (Lakoff, 2000, pp.159-160). This committee accepted Virginia Douglas’ (1972) reconceptualization of hyperkinesis as primarily a matter of inability to sustain attention for prolonged periods, whether due to motor hyperactivity or simple difficulty focusing. Earlier notions of ‘moral’ abnormality, of psychic conflict, and of imperceptible lesions within the brain were set aside, replaced with a definition encompassing a potentially much larger number of children. As opposed to acquired personality traits, it rested on Chess’ notion of ‘inherited temperament,’ and a promissory genetic-neurological etiology. Concretely, however, it was nothing more than a checklist. The DSM-II’s single sentence on a ‘hyperkinetic reaction of childhood,’ featuring “overactivity, restlessness, distractibility, and short attention span” (APA, 1968, p.50) was replaced in the third

edition with four pages of descriptive features and formal diagnostic criteria.¹³ Beyond these criteria from the DSM, there was an uncodified but general acceptance of the paradoxical subduing effect of psychostimulants as a diagnostic test (as described above and in Rafalovich, 2004), as well as a widespread adoption within diagnosis of the same Conners scale initially developed to test the efficacy of those same drugs (Conners, 1970; 1971). Henceforth a new psychiatric mainstream was established, its normative project an alliance of standardized questionnaires and pharmacology. On the surface Douglas' formulation comes full circle, citing the same passages from William James as did George Still, and conceiving the disorder of attention, likewise, as fundamentally a disorder of *will*. There were, however, several profound shifts marked by the DSM-III. Again, they make it difficult to regard Still's classification as in any way continuous with our own.

First, ADHD was viewed as potentially comorbid with 'mild or moderate mental retardation' (Webb, 1981, p. 52), thus rejecting Still's initial premise: that this was a disorder of 'moral' deficiency, one of conduct as opposed specifically to intellect. Conversely, the set of broadly delinquent, 'infracriminal' behaviours emphasized by the early medicalizers of 'moral degeneracy' in youth were simultaneously rearticulated into a distinct set of diagnostic categories, those of 'conduct disorder' and 'oppositional defiant disorder.' Skepticism toward these disorders likewise persists into the popular discourse of today, and is often lumped together with critiques of ADD/ADHD. These two shifts, along with the broader 'deinstitutionalization' of psychiatry and the diffusal of its authority into the community, removed the prior historical limits on the potentially diagnosable population. At the same time, the DSM-III explicitly disavowed the idea that deviant behaviours or

¹³ A child might be diagnosed with attention deficit disorder, after DSM-III, following affirmative answers by evaluators to two or three of the possible symptoms (lasting 'at least six months') under two or three of the subcategories, now separately defined, of inattention, impulsivity, and hyperactivity. Inattention included symptoms such as failure to complete tasks, and distractability; impulsivity covered acting before thinking, calling out in class, and excessive risk-taking; hyperactivity encompassed the traits of fidgetiness and behaving as if "driven by a motor" (APA, 1980).

social conflicts were to be seen as mental disorders, except when they represented a ‘symptom of dysfunction in the person’ (APA, 1980), an apparent limiting of its scope which underlines the replacement of external norms of conduct by internal psychobiological norms.

With this historical survey I have tried to highlight the crucial changes that produced a new “ecology of psychiatric knowledge” (Lakoff, 2000, p.157) wherein ADHD diagnosis could flourish. By privileging a biogenic conception over the psychogenic one which held sway even in the early pharmaceutical era, opening up to the medicalization of simple inattentiveness without excessive activity, and bolstering the apparent ‘scientificity’ of its claims through standardized questionnaires and randomized clinical trials, the psychiatry of the 1980s set the stage for an ongoing expansion of this diagnosis.¹⁴ In this ecosystem, as Lakoff puts it, “ADD thrived as a disorder that proved malleable through medication and which blamed none of the victims: neither parents, children, nor society” (2000, p.166). Thus there was relatively rapid mobilization of a community of ‘lay expertise’ (S. Epstein, 1998) around this appealing diagnosis, a coalition of support groups and advocacy organizations (such as CHADD, founded in 1987) which supported the consensus regarding medicalization and pushed for further legislative recognition and research funding (Lakoff, 2000, p.162). This is not to say that a countervailing discourse faded away; on the contrary, we can observe a continuing ‘turf battle’ around the margins of the psychiatric consensus.

My goal in this discussion has been not only to contextualize the development of the medical knowledge proper, and to furnish a background for assessing the contemporary coverage of the

¹⁴ The expansion of the category beyond those with hyperkinesis was indeed hesitant: the 1980 DSM-III includes both hyperactive and nonhyperactive forms of attention deficit disorder, while the 1987 DSM-III-R returns to the historical emphasis on hyperactivity. The revised manual allows only diagnoses of ‘Attention-deficit hyperactivity disorder’ and of ‘Undifferentiated Attention-deficit Disorder;’ regarding the latter it suggests that “research is necessary to determine if this is a valid diagnostic category and, if so, how it should be defined” (APA, 1987, p.95). The fourth and fifth revisions of the manual return to the DSM-III notion of a unified disorder with ‘predominantly inattentive’ and ‘predominantly hyperactive’ variants.

issue, but to support my contention regarding social construction: that ADD is neither discovered nor fictitious, but constructed in many layers across its history, from elements of observable reality, intermingled with heterogeneous assemblages of subjective, cultural, contextual, and interpretive material. The neo-Kraepelinian project of post-DSM-III American psychiatry has endeavoured to tame the flexibility of this interpretive material, and has no doubt dethroned psychoanalytic hermeneutics in much of the world, but *a priori* cannot wholly exclude interpretation from a study, in the end, *of* subjectivity. Prefigured by debates surrounding anti-psychiatry, drugs as tools of social control, and the medicalization of homosexuality, the DSM-III signals an era in which “psychiatry’s definition of the self becomes a public matter” (Healy, 2002, p.305). Within the public debate on this particular pathological entity, we can observe both an explosion in quantity of coverage, and an increasing diversity of views expressed, often echoing, resisting, or rearticulating positions from long ago in the scientific discourse. This chapter opens my description of ADD/ADHD, further developed in the next, as a ‘thing’ in the sense proposed by Latour: an historically-situated object of contested knowledge around which an assembly of different actors is gathered, a locus of both agreements and disagreements, disputation and consensus (Latour & Weibel, 2005, p. 5).

5.2 : Media research

Method.

I attempt to gain some purchase on the contemporary constitution of ADHD as a public, epistemic thing by examining coverage in the news media. In the remainder of this chapter, I detail the quantitative and qualitative findings of a content analysis carried out on a sample of recent news articles discussing the disorder. My aim is to analyse the quantity, tone, and themes of popular press coverage around the disorder, attending to four major areas of interest:

1. To what extent do mass media reports give voice to the view that increasing or excessive media use may cause attention deficits? What technologies are singled out in this regard, and how?
2. What are the attitudes expressed toward psychopharmaceutical treatment? How often is it discussed in comparison with more traditional psychosocial modalities of treatment?
3. What are the main issues of concern around rates of diagnosis and treatment for the condition? How are the concepts of over- and under-diagnosis presented?
4. What details, from what kinds of scientific studies, are presented? How do media stories construct their framing of the relationship between mind, brain, and illness – and how does this compare with framings within scientific discourse?

To this end, I collected a pilot sample of 100 news articles from 2003-2013 under the Lexis-Nexis subject heading ‘ATTENTION DEFICIT DISORDER,’ choosing from the thousand results within Major World Publications deemed most relevant by the site’s algorithms according to my own subjective evaluation of their relevance. Along with those news stories collected based on citations elsewhere in my secondary research, I read these articles both to collect qualitative material and to devise an emergent coding system whereby I could produce quantitative data. The core of my approach to media

coverage throughout this dissertation remains qualitative, and so I consider this a supplement from which limited, preliminary conclusions may be drawn, to be developed further by textual analysis.

All coding was performed solely by myself. While this is admittedly worrisome as judged against the ideal standard of quantitative rigour, which would demand assessment of inter-coder reliability, what follows is a qualitative analysis which employs quantitative coding as a guide. I follow the general outlines of a thematic analysis approach (Braun & Clarke, 2006), tracing out clusters of concepts based on my own situated knowledge and engagement with the data set. In writing this text, I worked by continually moving back and forth between the coding of my sample and the composition of my analysis. Rather than claiming in some loose sense that themes had emerged from my reading, I wished to measure their emergence; at the same time, measurement is not my primary focus here. Based on the initial pilot sampling, as well as prior research designs (cf. Racine et al., 2010; Schäfer, 2012), I constructed my main sample, and produced the outlines of a

Search queries used on Lexis-Nexis Academic Major World Newspapers:

1. BODY(television OR Internet OR technology OR computer*) AND SUBJECT(ATTENTION DEFICIT DISORDER 9*%) AND LENGTH >(500)) and Date(geq(5/8/1993))
2. BODY(television OR Internet OR technology OR computer*) AND SUBJECT(ATTENTION DEFICIT DISORDER 8*%) AND LENGTH >(500))

Figure 33 - Media sampling queries

coding structure which could reliably capture all major domains of interest within these articles. My sample would expand to encompass the past twenty years of coverage (1993-2013), delimited to those articles which could potentially discuss my core phenomenon of interest, the hypothesis of ADHD itself as ‘media effect.’ I wished to retain the global reach of the pilot sample, since there were interesting and relevant articles sourced from Asian and African publications which might allow some perspective on ADHD as not only a mediated but a globalized phenomenon.¹⁵ I elected, however, to exclude some of the niche publications found in my original sample (including specialist publications for psychologists and the pharmaceutical industry), thus I ultimately drew my sample from the Lexis-Nexis ‘Major World Newspapers’ collection, encompassing the major English-language ‘prestige’ papers from North America, Europe, Israel, Africa, Australia, and Asia, in addition to some more tabloid-leaning publications like the Mirror and Daily Mail from the UK, and the American Washington Times and USA Today.¹⁶

As with the pilot sample, my population of potential articles consisted of those under the Lexis-Nexis subject classification of Attention Deficit Disorder; however, by employing the percentage relevance scores assigned in the article metadata, I was able to limit those collected to the subset which discussed the disorder most extensively. Finding that those with a subject relevance

¹⁵ Schäfer (2012, p.656) notes in his meta-analysis of research on science and the media that coverage of non-Western nations has been sparse, with none of his sample investigating African media coverage and only 0.4% considering Asia. I will emphasize the limitations of my sampling from non-Western media sources – they constitute only a small percentage within my sample, and represent specifically an *Anglophone* demographic within their nations – but it seemed valuable to include them nonetheless. My conclusions will often be focused on the North American context, however, in some contexts I will argue that they are more widely generalizable, and in others that there are significant differences in the non-Western, non-Anglophone world.

¹⁶ Media research, particularly on science coverage, has typically focused on national ‘prestige’ press outlets with a reputation for journalistic quality. My sample includes these as well as papers which are more typically considered ‘tabloids,’ in an attempt to encompass a broader spectrum of the public discourse. Some valuable background studies specifically of tabloid journalism exist (Uribe & Gunter, 2004), as well as specifically of science coverage within them (Boykoff & Mansfield, 2008). I discuss both the quantitative and qualitative similarities and differences between different kinds of press below – though I interpret them differently in my context, I do concur that tabloid venues are more ‘steeped in opinion and commentary,’ presenting more substantial divergences from scientific consensus (Boykoff & Mansfield, 2008, p. 3).

<i>Technology and ADHD</i>	<i>Psychostimulant medication</i>	<i>Issues of diagnosis</i>	<i>Other</i>
Affirms television as potential cause	States proven efficacy	Claims of overdiagnosis	Claims of dietary causation
Affirms digital technology as potential cause	Addresses side effects	Claims of underdiagnosis	Claims of physiological environmental causation (<i>ie</i> , toxins)
Rejects technological causation	Expresses skepticism toward treatment	Encourages self-diagnosis	Discussion of pharmaceutical treatment
Discussion of technology neutral to questions of causation	Expresses concerns about diversion/misuse	Discusses gender imbalance	Discussion of psycho-social treatment
Discusses positive role for technology in diagnosis / treatment	Expresses pro-enhancement position	Claims diagnosis may be an asset	Discussion of disorder in children
		Neuroessentialism	Discussion of disorder in adulthood
		Anti-neuroessentialism	

score below 90% sometimes discussed ADD in a secondary way, and those below 80% tended only to mention the disorder in passing, I constructed my search queries (*Fig. 2*) to encompass this top 20% of articles relevant to the topic. My queries also excluded any articles of less than 500 words, since these were typically brief notices of new developments, often from a business-news perspective on the pharmaceutical industry, without significant discussion. The first query, encompassing those articles with greater than 90% subject headings relevance, returned 408 articles, while the second returned 190. All of the articles returned were downloaded, and the rough coding schema derived from my pilot sample was refined and finalized following a preliminary examination. Duplicate articles, letters to the editor, and false positives without a minimal level of discussion on ADD/ADHD were manually eliminated from the sample during coding. The entire corpus of articles was also subjected to frequency analysis using the Textanz software. Relevant high-frequency

keywords were selected based on subjective criteria of relevance to supplement my discussion of the coding results (see Figure 35).

Before commencing my coding, I considered some working hypotheses. Although I did not have specific advance expectations for the results on each element of my coding schema, I believed that:

1. There would be considerable discussion of the potential role of media in causing ADD/ADHD.
2. There would be majority support for the psychiatric consensus that stimulant drugs were a safe and effective treatment for this disorder.
3. Brain-based conceptions of the disorder, and discussion of the disorder in children, would be more prevalent in my sample.

Table 1 above presents the categories and groupings employed in coding these articles. These were largely established prior to the collection of my main sample, with the intent of capturing not only the main phenomenon of interest – the hypothesis of media effects – but also a range of associated elements prominent within the public discourse on psychology, neuroscience, mind, and brain. My aim in including these additional categories was to allow for a broader understanding of the public constitution of ADHD as an ‘epistemic thing,’ examining how the debates on media effects circulated within the same milieu as many other matters of concern. Foremost amongst these was the efficacy and desirability of psychostimulant treatment: how commonly did news media emphasize the pharmacological consensus that Ritalin and other stimulant drugs were proven to be safe and effective for the treatment of ADHD, and how commonly did they present potential side effects or skeptical views of this method of treatment? Another distinct set of concerns arising in

association with these drugs quickly became apparent, and so I added corresponding categories to my classification covering concerns over the diversion and non-prescribed use of these drugs, as well as the less common but intriguing argument that cognitive enhancement amongst ostensibly ‘normal’ individuals also ought to be deemed an appropriate use.

Neuroessentialism: “The causes of ADD/ADHD continue to be investigated but the current medical model is that it is a disorder born out of a deficiency in a key brain chemical called norepinephrine. Because norepinephrine is too low in those with ADD/ADHD, medications that stimulate the production of this chemical are often used.” – The National Post, Toronto (Brunshaw, 2011)

Anti-neuroessentialism: “In 1998, an overwhelming majority of experts attending the National Institutes of Health Consensus Conference, after days of reviewing all of the available evidence, agreed there is no compelling evidence that ADD or ADHD is caused by or significantly and reliably associated with physical or biochemical irregularities in the brain” – The Washington Times (Rosemond, 2004)

Figure 34 - Coding exemplars for neuroessentialism

The third set of categories encompasses matters of diagnosis. Do news articles more commonly raise the spectre of overdiagnosis and overtreatment with respect to ADD, or the opposite? Following my initial examination of the coverage, I was also interested in how frequently the gender imbalance in ADHD diagnosis would be discussed, and in claims that those with the diagnosis may find their disorder an ‘asset’ in some sense.

‘Neuroessentialism’ (Racine et al., 2005, 2010) was another discursive phenomenon which I sought to

capture. This posed some challenges for a reliable coding system, but employing the parameters laid out in Racine et. al (2010, p.728), I considered any implications that the ‘real’ nature of ADHD was to be found exclusively in the brain as falling within this category. Thus it encompassed claims that ADD involved differences of brain ‘chemistry’ or ‘wiring,’ as well as claims of specific neuroanatomical regions posited as responsible; the inverse category, ‘anti-neuroessentialism,’ was reserved for explicit arguments against the idea that the disorder was reducible to a matter of brain function. While common, instances in passing of neuro-essentialist tropes – wiring, chemical balance, and the like – were not coded within these categories, which I limited to causal claims

presented as authoritative (see Figure 34). The more common claims that the disorder involved some mixture of biological and social/environmental causality were not coded within either of these categories, nor were juxtapositions of these two sorts of claims, since either ‘balanced’ approach seemed to constitute a journalistic rejection of essentialism. Consequently, these two categories were considered mutually exclusive. The final axis of my coding structure includes three other sets of relevant issues not falling within the other areas. Two of these deal with the balance of coverage between the options for pharmaceutical and psycho-social treatment modalities, and between manifestations of ADD in childhood as opposed to in adulthood; the other covers two causal hypotheses which I quickly realized were common and relevant enough to include as separate coding categories, namely diet (especially in terms of additives) and exposure to environmental toxins.

In coding for my primary set of concerns, I separated out those articles which hypothesized a causal role for television from those which laid blame on newer digital technologies, including personal computers, electronic games, and the Internet. I also included categories for those which specifically rejected this account of media effects, and for those which simply mentioned one of my search terms without discussing their causal role. The last of these categories excludes the others within this axis; however, the same articles could potentially be coded under any or all of the remaining options. My aim was to separately tabulate all instances of these claims, even when more than one occurred in the same article, rather than attempting to decide which views were most prominent in a given piece. This strategy helps to capture the balance of coverage within articles as well as across all content. Having excluded ‘media’ itself as a term leading to an excess of irrelevant hits (due to its ubiquity in the corporate names of news publishers), I elected to include all those articles retrieved by my queries (*Fig. 2*), even where discussion of the technology’s potential causal

role in the disorder was absent. Rather than ruling these out of my sample as false positives, I retained them as a quasi-control set which would allow for better evaluation of the interplay between the media effects hypothesis and other salient aspects of the discourse. Television and digital technology are sufficiently omnipresent aspects of our everyday lives that they are regularly mentioned when not a major focus of news stories: mentions of ADHD sufferers working with computers, or finding support groups online, for instance, are quite common. Stories falling into this category were coded as ‘neutral’ mentions of technology. Their presence within my sample allows me to compare those articles which *do* discuss the media-as-cause hypothesis against the broader category of those whose content *might* contain such an examination (since they mention media technologies in some way) but do not. I found this valuable in assessing the ways that this particular claim about media effects fits into the broader ecology of mediated discourse on the disorder. Finally, I also became aware early in the coding process that a further category would be required in this grouping, for the relatively frequent articles extolling the potential positive value of technologies such as neuroimaging or neuro- and bio-feedback in diagnosing and treating the disorder.

Quantitative findings.

My findings regarding the hypothesis of ‘media effects’ were surprising in some respects, but confirmed my first general working hypothesis: that this would be a fairly common ‘angle’ for news stories about attention deficit disorder. (All results discussed in this section are listed in Table 2.) While I anticipated that computers, electronic gaming, and the Internet would be most commonly cited as cause for alarm, due principally to their novelty, I found that there was minimal difference in the rates of discussion for each. Television was affirmed slightly less frequently than digital technologies as a potential causative factor. As I consider in my qualitative discussion below, there

are some important distinctions to be made in the *way* that each medium is discussed. Some of the prevalence of coverage on television can be attributed directly to the 2004 publication of a study which argued for correlation between early exposure to television and the development of attentional problems (Christakis et al., 2004).

Only a very small percentage of articles explicitly rejected the idea that technology could be a causal factor in the rise of ADD diagnosis, in these cases typically as an anecdotal ‘foil’ of sorts to a mainstream biogenic conception: the disorder is caused *not* by ‘too much TV,’ but by an inheritable brain abnormality (e.g. Shapiro, 2007). The largest single category coded was of articles discussing technology in some fashion neutral to the question of causation, typically either in relation to careers or to public understanding of the disorder. Included in this category were articles where individuals were mentioned as having learned about the disorder via the Internet, as well as a surprising number of articles evincing concern over the ordering of pharmaceuticals online. While these were not precisely ‘neutral’ to other issues examined here, they did not take any stance on the media effects hypothesis. Although this was the most common individual category, it should be noted that of all articles mentioning technology and the media in conjunction with ADD, nearly as many offered causal claims involving television, digital technology, or both (121, or 35% of the total) as those which merely mentioned technology in a neutral way. The final category covered an unexpectedly large number of articles, most of which discussed the potential value of bio/neurofeedback or ‘brain

Table 2: Coding Results							
Expressed as percentages of total articles coded (N=346)							
<i>Technology and ADHD</i>		<i>Psychostimulant medication</i>		<i>Issues of diagnosis</i>		<i>Other</i>	
Affirms television as potential cause	22.3 (77)	States proven efficacy	18.2 (63)	Claims of overdiagnosis	28.3 (98)	Claims of dietary causation	16.2 (56)
Affirms digital technology as potential cause	25.7 (89)	Addresses side effects	16.8 (58)	Claims of underdiagnosis	17.1 (59)	Claims of physiological environmental causation (<i>i.e.</i> , toxins)	8.1 (28)
Total articles, either or both causal claims	35 (121)	Expresses skepticism toward treatment	37.6 (130)	Encourages self-diagnosis	8.1 (28)	Discussion of pharmaceutical treatment	60 (206)
Rejects technological causation	4.3 (15)	Expresses concerns about diversion/misuse	8.1 (28)	Discusses gender imbalance	6.7 (23)	Discussion of psycho-social treatment	55 (189)
Discussion of technology neutral to questions of causation	43.4 (150)	Expresses pro-enhancement position	4 (14)	Claims diagnosis may be an asset	12 (40)	Discussion of disorder in children	83 (288)
Discusses positive role for technology	20.5 (71)			Neuroessentialism	26 (90)	Discussion of disorder in adulthood	26 (90)
				Anti-neuroessentialism	4 (14)		

training' games in the treatment of ADHD, or the development of novel scanning technologies for diagnosis. Articles were only coded within this grouping if they discussed some specific, novel diagnostic technology; where technology only factored in as a vector for general information about the diagnostic classification of ADHD (*i.e.*, symptom checklists), the article would be grouped under the neutral heading. Again, these categories are not mutually exclusive.

The balance between causal claims regarding the media and neutral mentions of technology within my sample was somewhat surprising. I anticipated finding many mentions of technology which did not specifically address the notion that they might be responsible for the escalating diagnoses of the disorder. I underestimated how common these would be, however – it seems that technology is a sufficiently ubiquitous part of everyday life that it factors incidentally into all kinds of news stories. Not captured by the quantitative data, however, were the ways in which these non-causal claims suggested an underdetermination of the meanings attributed to media consumption. Anecdotally, many children diagnosed with attention deficits have no trouble focusing on highly-stimulating mediated activities like electronic games. This was a common refrain throughout my sample: as stated in one article which did offer both affirmative and skeptical views of media-as-cause, “her son could sit for hours concentrating on video games, it turned out, so she was certain there was nothing wrong with his attention span“ (Klass, 2011). This article was one of several which discussed specific research papers investigating correlations between media consumption and attention deficits. Like most such articles, it offered a relatively sophisticated take on causality, noting that increased ‘screen time’ may be interpreted as more an effect than a cause of attention deficits.

The majority of articles, however, were not focused on peer-reviewed research, and lacked such caveats. Thus observations like the one cited above were quite common, but often interpreted

quite differently. A high percentage of those articles advocating for media consumption as a causal factor in ADD pointed to intense media use by ostensibly disordered children as supporting evidence, while a great many of those within the ‘neutral’ category addressed the same patterns of use simply as reflective of the particular interests and aptitudes of the individuals in question. While I can draw no general conclusions based on the ratio of stories discussing the media effects hypothesis as a percentage of total issue coverage, my sample shows that roughly one third of the time technology is mentioned in an article on the disorder, that article does suggest it plays a causal role. The inclusion of the remainder, stories offering neutral or positive mentions of media, allows me to derive qualitative conclusions below regarding the intersections between claims of media effects and other areas of public concern, including skepticism or advocacy of pharmaceutical treatment and neuro-essentialism.

There were relatively few articles which specifically rejected the hypothesis of media effects, as defined in my coding. I expected this to be somewhat more common, given that this hypothesis seems to contradict the general biomedical consensus about the origins of this disorder. In the few articles which did take a negative stance toward media causation, this contradiction was evident: the claim of media effects served as a paradigmatic ‘anecdotal,’ popular epidemiology of ADD (P. Brown, 1987), contrasted in dismissive fashion with the ‘true’ neurobiological account of the disorder. This indicates both the prominence of the media effects hypothesis in popular consciousness, and the way such anecdotal claims in the press may serve as a fulcrum for *boundary-work*: the establishment of clear demarcations in the public sphere between scientific knowledge and non-scientific opinion (Gieryn, 1983; Moore & Stilgoe, 2009). Across the rest of my sample, however, there was a greater than expected tolerance for multifaceted, interactional accounts of the disorder, with many articles implying hybrid etiologies, blending neurobiological factors with

patterns of media use and other possible social and environmental influences. In this regard, one set of causal claims which I did not expect to find with such regularity in my sample were those relating to diet and to physiological environmental factors (that is to say material toxins rather than matters of social *milieu*). Although I was aware of the Feingold food-additive claims (Feingold, 1975; Smith, 2011), it was surprising to find them almost as prevalent within this sample as the main technology-oriented claims of interest, found within 16.2% of my articles, as compared with 22.3 and 25.7% for causal claims regarding television and digital technology, respectively. Many articles combined all of these claims, as part of a broader critique of medicalization, which I discuss below in greater detail. Others simply included brief claims about dietary causation, in a gesture of balance with more mainstream biomedical accounts, while a small handful discussed specific scientific investigations of the relationship between behaviour and food additives (e.g. McCann et al., 2007).

This was one context in which there were substantial differences in coverage between regions: by far the highest percentage of articles discussing this issue were in newspapers from the United Kingdom and Europe (see Table 4). Whereas 25 percent of articles from this region addressed the claim of dietary causation, only 7 and 8 percent, respectively, of articles from the United

Figure 35 - Prominent keywords in corpus (with total occurrences)

‘Children’ (3253), ‘child’ (1292)
 ‘School’ (1615), ‘teacher’ (447)
 ‘Parents’ (1221) – ‘adults’ (439)
 ‘Drugs’ (1050), ‘Ritalin’ (985),
 medication (569) – ‘therapy’ (199)
 ‘Brain’ (908) – ‘mind’ (202)
 ‘Television’ & ‘TV’ (683)
 ‘Computer’ (328), ‘Internet’ (273)

States and Canada did so. Although caution should be exercised in accepting the broad rhetoric of a ‘precautionary principle’ in European policy (Sunstein, 2005, pp. 20–23; J. B. Wiener & Rogers, 2002), on the specific issue of synthetic food additives and dyes, the region manifests stronger public concern and a more restrictive regulatory regime. The greater prevalence of this coverage in Europe may be an artifact of my sample, but as there is no reason to suppose that delimiting my articles to

those mentioning media technology would differentially retrieve more mentions of diet in Europe as opposed to in North America, I believe this accurately reflects patterns of coverage.¹

Predictably, I found much more discussion of the disorder as manifested in children, as opposed to in adults. ‘Children’ is itself the single most frequent word in the whole-sample corpus, with 3,253 occurrences – as well as 1,292 of ‘child’ and 1,221 of ‘parents’ – compared with 439 occurrences of ‘adults.’ This implies that children remain, by a large margin, the primary focus of concern regarding this disorder. My sample suggests that the medicalization of children’s behaviour, and treatment of children with psychopharmaceuticals, is widely seen as more worrisome than medical interventions for adults. In media discussions of ADHD, adults are more likely to factor in as caregivers – parents or teachers – than as subjects of diagnosis in their own right. This is of course understandable in the context of this disorder’s history: it has always been conceptualized primarily as a disorder of children, the symptoms of which typically would either dissipate or transform into more serious patterns of delinquency by adulthood. With the shift away from a conception of the disorder as one of ‘moral control,’ and its uncoupling from notions of criminal-mindedness, the doors were opened to diagnosing adult difficulties in sustaining attention by the same category. The articles within my sample track the gradual increase in public awareness and acceptance of this new category, illustrating the claims and contestations involved in the construction of a new species of personhood: the adult with ADD.

¹ Beyond my sample there is nevertheless some coverage of the relationship between food additives and ADHD within North American media; such stories often focus on the greater restrictiveness of EU law as opposed to the American FDA and Canadian Food Inspection Agency, as well as more specifically on the question of diet, whereas European news (in my sample predominantly from England) seems more likely to mention dietary factors as general background to stories about the disorder, regardless of their specific focus. This also seems to support the view that diet is a more salient matter of concern in European public discourse than in the United States.

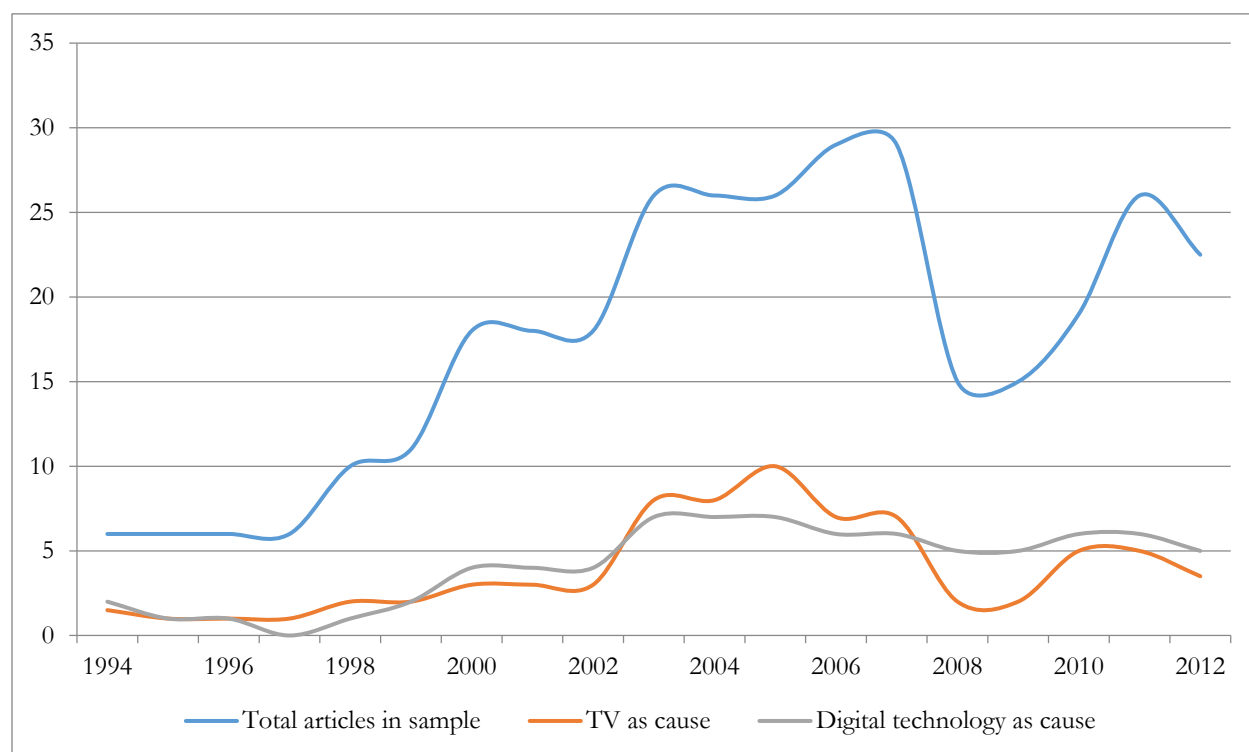


Figure 36 - Breakdown of sample by year

Over the period investigated I found a significant, though uneven increase in coverage of ADD/ADHD generally, as well as of stories specifically mentioning technology. Figure 36 graphs the breakdown of articles with three-median smoothing applied to the data. There is a peak in coverage of causal claims around media subsequent to the publication of Christakis' television study in 2004 (Christakis et al., 2004), and of a study with similar conclusions in 2010 (Swing, Gentile, Anderson, & Walsh, 2010), as well as a peak in 2006 unassociated with any specific research findings. The increase in coverage within my sample is not as significant or unequivocal as the increase which precedes my sampling period, that is to say from the enshrinement of ADHD in the 1980 DSM-III up until 1992. My sample was limited to the prior twenty years at the time of sampling (1993-2013) not simply out of convenience, but because there was minimal coverage of the

disorder prior to 1993. Despite the long pre-history presented above, and a significant controversy over drug treatment for hyperactivity in the 1970s, there was little discussion of the newly-christened disorder in the popular press of the 1980s. Using my Lexis-Nexis Major World Newspaper queries for the past twenty years (Figure 33) returned 508 results; using the same queries for the date range 1980-1993 retrieves only four. Even without limiting the sample to articles mentioning media technology, searching for all articles with greater than 80% subject relevance for attention deficit disorder from 1980-1993 returns only 48 articles from the entire 13-year period. By way of comparison, just the period from June 2012 - June 2013 returns 126. This implies that the disorder itself, alongside the potential links – both negative and positive – with media technology, have both been objects of sharply increasing awareness and concern since 1990, correlative with the ongoing expansion in diagnosis and treatment. Medicalization of behaviour is more sharply contested, of course, when diagnostic categories expand to encompass more individuals, particularly children.

The case of ADHD also suggests, however, that popular media coverage is strongly influenced by pre-existing moral framings of the behaviour in question, and also of the treatments available. Previous studies of media coverage on autism, the other most prominent ‘epidemic’ of mental disorder among youth, indicate a wide range of common themes and foci of discussion but very few claims that it constitutes a false pathologization of mere normal variation, even where overdiagnosis was sometimes flagged as a concern (Clarke, 2012; S. C. Jones & Harwood, 2009; Kang, 2013). In my sample, by contrast, overdiagnosis was a common concern – discussed in 28 percent of my articles, far more than underdiagnosis – and often linked up with broader skepticism about the reality of the disorder itself. Autistic symptoms, particularly at the more severe end of the spectrum, would likely once have been flagged as forms of mental retardation (Eyal, 2010). Concerns of overdiagnosis in that case bear more on the least-severe end of the spectrum, where it

encompasses deficits in social interaction which may once have been seen as simple eccentricity. Abnormality, in the case of autism, seems to be more readily and less controversially medicalized. In the case of attention deficit, however, the entire spectrum of problematic behaviours, from hyperactivity to distractedness, have been framed since time immemorial as problems of parental discipline and moral responsibility, not of intrinsic abnormality. This represents a powerful competing discourse in the public ‘turf battles’ over the medicalization of ADHD. In stories on autism, ‘social responsibility’ typically figures in as the responsibility of society to improve treatment and fund research into the condition (Kang, 2013); just as often, in stories on ADHD, medical diagnosis and treatment are themselves rendered as abrogations of the social and moral responsibilities which have long defined childrearing.

The other crucial factor is the symbolic framing of treatment options. In the era before Charles Bradley’s Benzedrine studies, when treating abnormal children meant simply a more refined and targeted mode of pedagogy, it would be hard to imagine such controversy over medicalization. Psychiatric medicine targeted a different set of techniques on a smaller set of subjects. Contemporary discourse on ADHD, however, is shaped by the history of cultural meanings attributed to psychostimulant usage, by the recollection of an ‘amphetamine epidemic’ driven by doctors and the pharmaceutical industry, and by the chemical similarities between the major drug treatments and highly-stigmatized ‘recreational’ drugs (particularly cocaine and methamphetamine). News reports claim for instance that Ritalin is “uncannily like cocaine in terms of what it does to the brain,” highlighting its diversion from prescribed use for recreational purposes with slang terms like ‘kiddie cocaine’ and anecdotes of teenagers insufflating crushed tablets for a stimulant ‘high’ (Snow, 2010). The widespread skepticism toward psychopharmaceutical treatment was a surprising finding of my research: present in 37.6% of my articles, substantially more than the 18 percent which stated

the consensus view as derived from clinical trials. I had expected some skepticism, certainly, but not a greater percentage than those articles which emphasized findings of safety and efficacy. The recurring allusions to illicit, recreational drug use in these articles suggest that preexisting stigmatization of psychostimulants represents a major driver of popular dissent against the pharmacological mainstream.

Many articles did present both kinds of claims, in a standard journalistic gesture of ‘balance,’ contrasting the claims of skeptics with those of mainstream psychiatrists who affirmed the efficacy of Ritalin and other psychostimulants. This was unexpectedly rare, however: of the 206 total articles which discussed pharmaceutical treatment, 130 offered skeptical views and 63 presented scientific claims of efficacy, while only 26 presented both. Only fifteen articles discussed side effects in conjunction with claims of pharmacological efficacy. Some science-journalism commentators have decried this pattern of coverage as classic ‘scaremongering’ and controversy-stoking behaviour from the popular press (Raeburn, 2010). I find this critique not especially compelling. Of course controversial topics are likely to garner more media coverage – this does not imply that the controversy itself is a baseless figment of the journalistic imagination. On the contrary, it seems misguided to take for granted both the validity of the psychopharmaceutical consensus and the marginality of its critics. There is a certain irony in targeting the *New York Times* for a supposed publication bias against psychostimulants, when trenchant critiques of the pharmaceutical industry rest on its own publication biases (Angell, 2004; Goldacre, 2012).

I am more troubled by a rather bifurcated media discourse: contrary to the typical expectation of journalistic balance, articles offering skeptical or affirmative views on pharmaceutical treatment are more likely *not* to present alternatives. There seem to be two worlds of coverage. One adheres more closely to apparent norms of ‘prestige’ science journalism, presenting basic and

medical research in simplified form, and then another is dominated by skeptical opinion columnists bemoaning the drugging of children and the eclipse of old-fashioned disciplinary morality. Serious, balanced coverage which attempts to do justice to the complexities of both mainstream and dissenting opinion was almost nonexistent in my sample (and not unexpectedly, in spite of the criticisms cited above, this was found most often in major prestige news outlets like the *NYT*, *WSJ*, and *Washington Post*). Some of these stories – as well as some of the most troubling – are discussed in greater detail within the following section. Qualitative analysis reveals that even where the same article includes both perspectives, the effect may be undesirable, occluding crucial differences between the kinds of networks which individual spokespeople represent.

Another salient finding was a limited but recurrent discussion of the gender imbalance in diagnosis, with 6.7% overall discussing this issue. The full spectrum of responses were in evidence, with some articles presenting it as uncontroversial that more boys simply have the disorder, while others claimed this as evidence for overdiagnosis in boys or underdiagnosis in girls. Beyond this important area of concern, a wide range of gender issues emerged in these articles. As I expand upon below, many of these articles bore out Ilana Singh's observation that childhood behavioural disorders regularly serve as pretexts for 'mother-blaming' discourse (Singh, 2004). This was particularly evident in the British tabloid press and in conservative-leaning opinion columns. Another significant regional difference was found in the balance of coverage in Asian newspapers, most striking on issues of diagnosis. Where newspapers from the United States, Europe, and Australia presented claims of overdiagnosis in an average 31% of their articles, these were present in only 3% of articles from Asian publications. By contrast, the Asian press commonly presented

underdiagnosis (typically by comparison with American and European prevalence rates) as a major issue of concern, and regularly encouraged self-diagnosis or parental diagnosis (see Table 4).

In keeping with discussions of the globalization of Western mental health (Ballenger et al., 2001; Kirmayer, 2006; Watters, 2011), this suggests that popular media discourse in Asia is at present especially favourable to the medicalization and pharmaceutical treatment of attention. While I cannot draw conclusions with any certainty, this is likely due to a combination of factors, including the relative unfamiliarity and rarity of the diagnosis in the region, an editorial tone of technoscientific optimism in the papers sampled, and, most controversially, the efforts of pharmaceutical companies to develop a very lucrative new market. This is a valid concern, and no doubt some of the physicians consulted in these stories have professional connections with the industry, yet my data does not indicate any kind of coordinated conspiracy in this regard: though Asian coverage is the most concerned about underdiagnosis, and presents far fewer skeptical views of medicalization than the European or North American press, there is also far less discussion of pharmaceutical treatment than of psychosocial modalities. Further research is needed to investigate these differences between patterns of coverage in Asia and in the West, and to consider the role played by popular media in the globalization of mental health and psychopharmacology.

Table 4: Breakdown by region
(as percentages of total articles published in region)

Region	# of articles	Media technology					Psychopharmaceuticals					Diagnosis						Other						
		TV as potential cause	Digital tech as potential cause	Rejects tech. causation	Neutral on technology	Positive uses of technology	Efficacy stated as proven	Side effects discussed	Skepticism	Diversion/misuse	Pro-enhancement	Overdiagnosis	Underdiagnosis	Self-diagnosis	Gender imbalance	May be an asset	Neuroessentialism	Anti-essentialism	Dietary causation	Environmental cause	Pharmaceutical treatment	Psychosocial treatment	Adult	Child
United States	75	24	33	5	35	24	21	13	32	5	1	27	13	5	7	11	27	8	7	3	63	63	28	80
Canada	37	8	19	3	62	16	16	16	30	5	3	16	22	5	5	32	16	3	8	8	57	59	41	70
Europe/UK	119	30	29	3	43	13	14	18	45	10	8	35	16	3	6	9	21	3	25	11	62	49	24	80
Australia/NZ	73	16	19	5	45	26	26	18	45	7	1	32	16	7	8	11	30	5	15	8	64	49	30	90
Asia	32	16	16	6	38	34	9	13	6	6	6	3	22	13	6	3	41	0	13	6	31	63	6	97
Other	10	30	40	0	50	10	20	30	60	30	0	60	30	0	10	0	40	0	30	20	70	60	10	100

Qualitative analysis.

The portrayal of attention deficit disorder in the media is of special relevance to STS, not only as a case study in the public understanding of medical science and pharmaceutical technologies of behaviour, but as the site of public debate on longstanding concerns of the field: the meaning of social construction, and, more profoundly, the causal interconnections between minds, societies, brains, and technologies. I wish to engage with these compelling broader questions through this media research, as much as the specific hypothesis that the media may cause attention disorders. Through these debates, the media are recursively both actors in and forums for discussion of social construction. Typically implicit, there were some instances where these questions were dealt with directly. One article first cites claims that Einstein and Mozart were “historical greats” who grappled with attention deficit disorder, notions held by ‘advocates’ which it opposes to the arguments of historian Matthew Smith. He adopts a stance on the disorder’s social construction similar to my own (Smith, 2012). Smith’s claims are reported following a Congress of the Humanities and Social Sciences in Ottawa, Canada, in a not grossly distorted but brief, simplified, and controversial fashion:

Medical historian Matthew Smith contends ADHD did not exist before the 1950s and says the disorder emerged only because of the social, economic and cultural changes of the last half-century. Hyperactive behaviour only became a problem - and one that required treatment - when children spent less time in farm fields and more time in classrooms, says Smith . . . ADHD is the most prevalent mental health problem among children, affecting from 5 to 12 per cent of kids. Whether ADHD is a disorder needing treatment depends upon society’s expectations and a person’s environment, says Smith . . . “It is a real disorder,” he says, “but one that is represented by a series of symptoms that are only problematic in a certain set of circumstances.” (Ogilvie, 2009)

Here we find a stark framing of social construction. The article captures Smith’s core contention, that ADHD has not ‘always existed’ waiting to be discovered, but constitutes rather a set of behaviours which became targets of medical intervention in a particular historical context, and were grouped together by psychiatry in a shifting, contingent fashion.

The journalist faithfully reports that Smith does not view his research as debunking the reality of the disorder, but renders it in terms apt to suggest otherwise, evidently provocative to those involved in treating the disorder (while also inserting a seemingly out of place statement of fact about its prevalence, without attribution). The article goes on to furnish some predictable responses:

Clinicians, psychiatrists and experts in ADHD are quick to dismiss the idea that hyperactive kids only became a problem in the super-structured society of the 1950s...ADHD is one of the most credible and treatable psychiatric disorders, a fact that cannot be challenged by historical research, says Dr. Atilla Turgay, director of the Scarborough Hospital's attention deficit hyperactivity disorder clinic. Studies have shown people with ADHD have smaller, lighter brains than the general population, for example, and their brains have metabolic abnormalities, including poor glucose consumption, Turgay says. . . . Smith is not the first researcher to claim ADHD is a social construct, an argument Turgay says the profession "has already dealt with and discarded." And, Turgay adds, dismissing ADHD as a "real" disorder is dangerous since it might prevent people from seeking treatment. (*ibid.*)

In this miniature recapitulation of the so-called 'science wars,' we find a paradigmatic mediated presentation of boundary-work, as medical spokespeople enroll scientific and moral propositions to denounce the encroachment of this historian into their sphere of authority. Citing general evidence of organic differences between populations diagnosed with ADHD, Turgay suggests that decreased total brain volume and reduced glucose consumption mean that this population must have in fact existed throughout history, waiting to be discovered, diagnosed, and treated; an analysis in terms of social construction is presented not simply as a threat to internalist historical accounts, but normatively dangerous for its potential role in dissuading the afflicted from seeking treatment. Where critiques call attention to the constructedness of psychiatric diagnoses, such responses are common.

The above article is typical in that it presents two sharply opposed viewpoints in a concise fashion – the prototypical 'fair and balanced' pattern of coverage – without engaging substantively with the issues each participant raises in relation to the other, or with the evidence for claims presented as scientific. It fails to address, for instance, that as imaging researchers are careful to

point out, “anatomical differences are detected when looking at *groups* of children with and without ADHD. An individual may or may not have alterations in these brain areas,” hence ‘normal’ brain scans should not deter, nor those resembling the findings in diagnosed populations indicate treatment (Giedd, Blumenthal, Molloy, & Castellanos, 2001, p. 45). Nor does it explain that blood glucose consumption in the brain is the very currency of functional imaging research, as I have discussed in previous sections. Hence without further details on areas of the brain said to have ‘poor’ consumption or the controls against which they are compared, Turgay’s claim has little substance. More fundamentally, however, the presence of real and demonstrable differences in the brains of those diagnosed with ADHD does not contradict Smith’s claim of social construction. The point of emphasizing its historical particularity as a medical diagnosis is not to undermine but contextualize the reality of the disorder: to articulate how, even if we imagine a mature neuroscience of ADHD, with a full account of its distinctive markers in the brain, the scientific and clinical enterprise remains pervasively social. A longer piece may have allowed for more extensive dialogue between these two positions, but this kind of brief and dichotomous framing proved far more common in my sample.

Such framings ignore the true import of historical and social accounts of the disorder, irrespective of any evidence for its organic causation. There are important questions to raise about the validity of the scientific claims in this article and many others, but even a hypothetically irreproachable brain-based account leaves the origins of ADHD doubly constructed: in and by medicine, as social factors lead a set of behaviours which might in other contexts have been regarded as within the range of normal or perhaps ‘sinful’ conduct to be defined as pathological and in need of medical treatment; but also in the plastic medium of the brain itself, where behaviour is shaped only in part by innate or inherited characteristics, and far more by ongoing ‘rewiring’ through

social interaction. Between these two poles of influence on the brain and mind, we find the major contemporary iteration of the ‘nature-nurture’ debate, and the focus of some polemic between natural scientists and ‘humanists,’ broadly defined (e.g. Pinker, 2002; Sahlins, 1976). My stance throughout is that both nature and culture – along with external media – are integral to the emergence of mind from brain, in a ‘looping’ and reciprocal interaction, with the extent of each influence to be progressively better understood by diverse modes of research. Hence without allying myself to particular claims about human nature or cultural effects, in describing this range of controversies, I find certain media accounts can be normatively assessed as poor when offering an oversimplified narrative running in either direction – be it that ADHD is a genetically-determined biological disorder ‘like any other,’ or a fictitious construction deployed for profit and social control.

Like any patterns of behaviour, attention deficit disorders are sure to be correlated with significant general differences in the brain, but nevertheless the meanings attributed to and the mental experiences associated with them remain individual, contingent, and social. Like the vast majority of those in my sample, the above-cited article deals with none of these complexities. Though several other articles presented claims in a fashion which obscured their authors’ credentials to dubious effect – for instance citing criticisms of psychiatry by the Scientology-affiliated Citizens’ Commission on Human Rights without further background on the organization (Womack, 2006) – the effect of the credentials provided in this piece is quite the opposite. The framing undermines its superficial balance, as it contrasts the controversial claims of a graduate student in history with responses representing the ranks of clinical ‘experts’ in ADHD, and a professionalized mainstream which has ‘dealt with and discarded’ claims like Smith’s.

Where social construction is discussed explicitly in the media, the treatment is often similarly dubious. However, there are many telling examples in which the same debates are indirectly given

popular form, often in relation to the question of diagnosis. Is the ‘epidemic’ of ADHD based on the discovery of a biologically real, previously underacknowledged disease? Or is it in some way a false construction? This question stands at the centre of numerous articles from around the world, often presented in this fashion even where the article argues for not only its reality but its underdiagnosis. The *Glasgow Herald*, for instance, argues:

The popular belief is that increasing numbers of children are being diagnosed with attention deficit disorders as an excuse for bad behaviour and are prescribed medication to achieve docility by chemical means. That the problem turns out to be quite the reverse, a vast under-diagnosis, is alarming... There could be up to 34,000 children and young people in Scotland with undiagnosed ADHD. (“Children with ADHD; Apparent vast under-diagnosis is alarming,” 2008)

The claim of ‘popular belief’ is hard to quantify, but is no doubt based in part upon a substantial number of articles from the British tabloid press which make claims to that effect. Many blame television and video games, while also implicating overly restrictive, demanding schools, and – binding it all together – poor, neglectful parenting. These are even occasionally tenuously linked up with brain-based accounts, showing that while there is some tension between the two, such conservative positions are by no means mutually exclusive with them:

what’s happening is that some people carelessly have children, don’t look after them properly and then some-one else - the state, the care system - takes over. The process of abandonment combined with lack of love and maltreatment is actually altering the young brains of these children - causing brain damage, really. That’s how you get disorders like ADHD: the odd behaviours point to something having changed in the chemistry of the brain, perhaps even in its structure. (A. Brown, 2011)

Here the narrative of organic causation in the brain is presented not as contradicting, but complementing an account based in parental irresponsibility: ‘brain blame’ allied with parent-blame (cf. Singh, 2004). Excesses of ‘screen time’ are now a ubiquitous peg for this argument. In one unexpected variant of the claim that media exerted effects on ADHD diagnosis, others who claimed overdiagnosis placed blame not only on parents, teachers, and doctors but specifically on the use of the Internet for ‘lay’ diagnosis (Roberts, 2006). These claims, and the political implications which

often accompany them, may be worrisome, but in fact their picture of the interaction between brain and culture may find as much if not more support in scientific research than a straightforward narrative of brain illness as bio-genetic destiny.

Other articles which featured more extensive discussions with clinicians, however, offered a more balanced and nuanced take on the uncertainties of diagnosis. One manifests quite clearly a *labeling paradox* at the core of ADHD treatment: in the same piece, a mother states of her son that “for me, it’s like, label him with something so we can help him” while, further down, a psychiatrist refers to the practice of prescribing drugs for purposes which have not been formally approved or tested, saying that “everything we do in child psychiatry is off-label... You are confronted in the office with a dilemma. Okay, we don’t have a lot of research, but I have a kid who’s really struggling here” (Mick, 2008). ADHD is an exception in this regard, in that since the origins of psychostimulant therapy, the subjects of research have been predominantly children; the long-term effects of these drugs beyond the immediate subduing of behaviour remain under-researched, however, and many other widely-prescribed psychiatric medications have not been extensively tested on children at all. Some of the most popular stimulants, like Adderall, presently carry ‘black-box warnings’ upon their labels, indicating serious potential side effects. This drug specifically became a major focus in the Canadian press after it was implicated in several deaths by cardiac arrest, leading to the suspension of its national approval in 2005, which was eventually reversed on condition that the drug’s label be revised to indicate the increased potential for ‘adverse cardiac events’ (Levine, Gow, & Shea, 2005).

Diagnostic and pharmaceutical labeling are key ‘attractors’ in the looping trajectory of ADHD – and of pharmaceuticalized ‘patienthood’ more generally. The well-known ‘butterfly’ graphs of Lorenz attractors, somewhat *cliché* and recently discredited as a mathematization of

‘positivity ratios’ in psychology (by a collaboration including the infamous Alan Sokal: N. J. L. Brown, Sokal, & Friedman, 2013), nevertheless offer a fruitful analogy. Since Bradley’s studies, psychostimulants and their effects on disruptive, hyperactive children have been at the centre of a chaotic swirl of claims, eventually coalescing more tightly around the established DSM-III diagnosis of ADD/ADHD and the drug labels which indicate regulatory approval to treat and to market for that particular condition. The latter role of a drug label is considerably more important than the former: as the psychiatrist cited above mentions, clinicians are always free to prescribe ‘off-label’ if they believe their patients will benefit; the significance of the label is that manufacturers cannot specifically recommend drugs for anything but the approved conditions, and some of the largest legal penalties levied against drug manufacturers have been for precisely this offence (most notably, in pediatric psychiatry, for Johnson and Johnson’s efforts to encourage prescribing its Risperdal for the controversial unapproved condition of pediatric bipolar disorder).




Figure 37 - Lorenz attractor

If we imagine the labels of drug and disorder as the twin foci of this strange attractor, circulating around them are not only the professionals who define them but the individuals who are subject to them, in an iterative, mutually constitutive two-way traffic. Clinicians make decisions about treatment and diagnosis based on the two labels; they act upon the patients’ self-conception and social position as the drug treatments simultaneously act

upon their physiology; the patients’ individual responses and the clinicians’ replies determine the course of treatment (or non-compliance and resistance); all of these in aggregate, through

professional channels and popular discourse, combine to shape the constitution of the disorder as a clinical object and as a public thing. Unlike the models of Lorenz, however, the nature of the attractors cannot be specified in advance or in the abstract. The labels are just our reigning approximations, the ‘boundary objects’ (Star, 2010) which presently exert their power by satisfying the informational demands of many heterogeneous groups – scientific researchers, clinical psychiatrists, clinics, insurers, patients, and public spokespeople. Yet this is an uneasy and incomplete satisfaction, and these approximations are in numerous cases explicitly resisted. I invoke the trope of the strange attractor to emphasize my conviction that there *is* an underlying reality shaping (but not necessarily determining) the becoming of this thing labeled ADD, as well as of the drugs labeled as marketable for its treatment. I also believe that our knowledge can in some nonlinear and irregular fashion approach this reality more closely. But the reality is not to be equated with the labels, nor has it been revealed through some direct line of progressive discovery beginning with a German fable and culminating with the DSM-III.

Considering one of the drugs themselves offers the best illustration of my point here, as the half of the equation where most would agree that science has the better ‘objective’ knowledge. The amphetamine molecule itself is well understood. Chemists can represent its complete structure in a standardized form (*(R,S)*-1-phenyl-2-aminopropane), and corporations are able to manufacture on demand as much of both its *levo*- and *dextro*- isomers as they are legally permitted by their governments, as well as combine it with other compounds in reliable ways for different desired properties. Yet the molecule itself does not constitute *the drug*, as such. This demands consideration of its pharmacology, hence its interaction with the body. Here, knowledge of drug effects in the body may consist of clinical observations, like Bradley’s and many others, or studies in animal models and human subjects of the drugs’ physiological effects in the brain and central nervous

system. Considerable research has been done in this regard, and many findings are generally accepted: for instance that amphetamines promote release of dopamine and serotonin in various parts of the brain, in many of the same regions as other potentially addictive stimulants like cocaine, and related pharmaceutical compounds like methylphenidate. Likewise, it is well understood that different routes of administration sharply alter perceptions of the drug, with nasal insufflation and intravenous usage having more rapid bioavailability and accompanying subjective euphoria. Far from exhausting the reality of the drug, much like elsewhere in science, this research is an iterative process linking together a diverse range of techniques and actors through institutions, inferences, analogies, models, and translations: from rats and human patients to brain scanners and simulated neural networks. Further, it seems definitively beyond the grasp of objectivity to comprehend the full spectrum of individual, contextual responses to a compound, on a physiological, far less an experiential level. And ultimately it is by the compound's use in society that it becomes a *drug*, with all the cycles of institutional regulation and individual inter-reaction which accompany any truly *popular* drug. The drug label, though treated as definitive in some contexts, is but the focal point of many circulating factors in the popular shaping of psychostimulant use.

The same loop of objects, uncertainties, and translations holds on the side of the clinical diagnosis, with the further complication that labeled subjects respond, in often unpredictable or resistant ways. Both labels occasion reflection and decision on the part of various actors. Will my patient be insured for, or I be insulated against malpractice lawsuits if I label them with this condition and prescribe that drug? What resources will become available to me if I, or my child, or my student receives this label - a commonly cited concern in media narratives, which also incidentally drove the aforementioned historian to his research on ADHD (Smith, 2012). How does being diagnosed with this disorder make me feel, and others feel about me? These are difficult

questions, fraught with ambiguities and uncertainties in every particular case. The conclusion should be that diagnostic labels are neither fictions nor reflections of a pure objective reality, but strategic and performative constructions, constrained only loosely by observable phenomena while enacting constraints of their own. Across the personal narratives presented in the media, they come to be seen as valuable in certain individual contexts, as an internalized and positive dimension of self-understanding, and in others as repressive and negative. And in most cases, the most salient aspect of the labeling has to do with the fact that it includes access to stimulant drugs, often where the patient is a minor not considered legally competent to make their own treatment decisions, and some combination of parents, physicians, or teachers is encouraging drug therapy.

In popular narratives, skepticism of drug therapy was regularly presented in contrast to mainstream medical claims of efficacy, but the former quantitatively and qualitatively outweighed the latter. The more controversial claims about, as one regular press source calls it, the ‘multi-million dollar medical myth’ routinely gathered coverage and headlines, while statements about the effectiveness or pharmacological value of the drugs were relegated to the sidelines. Concerns seemed to be divided neatly according to age, with those regarding children far more common, having to do with ‘drugging’ (a recurrent phrase with definite negative connotations) for behavioural control, while in the context of older adolescents and adults, concerns are voiced more about the diversion of pharmaceuticals for recreational or ‘enhancement’ purposes. These columns were a site for the renegotiation of what had seemed to be a settled conclusion within American psychiatry: that the drugs worked to treat an internal dysfunction ‘in the person’ rather than operating to stimulate or calm behaviour in accordance with external norms of conduct (Comstock, 2011). This is by no means evenly accepted in popular media discourse. There are occasional claims adhering closely to this view, that for instance norepinephrine or dopamine are ‘too low’ in those with the disorder, and

thus medications that ‘stimulate the production’ of these chemicals are used to treat that deficiency (Brunshaw, 2011). Evidence does exist for differences in the dopamine systems of populations diagnosed with ADHD, but any full neurobiological account of the disorder is unlikely to be as straightforward as a deficiency in a single ‘key brain chemical’ (*ibid.*).

Rather than such specific neuroscientific framings, many articles opt for metaphorical characterizations of the effects of treatment. Here I found a sharp contrast between the kinds of metaphors employed by mainstream researchers and clinicians, and those who adopted more skeptical positions. It is inevitably difficult to characterize the ‘slant’ of an article, which I deliberately avoided trying to quantify by simply counting occurrences of different claims. Yet particularly with opinion pieces there were some which unquestionably sided with the medicalized conception of ADD, and others which adopted a sharply opposing stance. In fact, among opinion pieces, the latter position was more common, often taking shape as a conservative critique of psychiatry for its supposed role in the abnegation of moral responsibility. In journalistic articles, spokespeople for both points of view were more likely to be selected, with any apparent slant typically leaning more toward medical authority. The claims were accompanied throughout by two differing sets of technological metaphors. Some drew analogies with long-accepted medical interventions whose value is taken as uncontroversial. Occasionally there were references to insulin or antibiotics, but the most telling and common of these tropes referred to *eyeglasses*: parents with hyperactive or inattentive children, “most specialists agree, should be no more reluctant to try Ritalin than to give eyeglasses to a nearsighted child” (Allen, 2001). Rendered in even stronger neurobiological terms, this trope also animates a children’s book about hyperactivity called *My Brain Needs Glasses* (Vincent & Hoff, 2010). This metaphor was found in a few comparable contexts within my sample, more widely in online and broadcast discourse, and in professional literature as well,

typically coupled in this fashion with the claims of a professional consensus. One instance in an education journal, widely repeated online, couples the eyeglass and insulin analogies together with a contention that ‘addiction’ to psychostimulants was not a major worry, and children on the drugs may be “dependent,” but only “in the same sense that a person with diabetes is dependent on insulin or a nearsighted person on eyeglasses” (Prater & Pancheri, 1999). Such analogies are the currency of understanding and persuasion with respect to new technologies, as the novel and mysterious are interpreted through the lens of the old and familiar. Occasionally proponents of drug treatment will simply deny the possibility of tolerance or dependence, but with amphetamine compounds long stigmatized as drugs of abuse and addiction, this strategy appears less credible. Instead, analogies with uncontroversial medical technologies on which patients also depend for life serve to symbolically domesticate long-term stimulant treatment, even where it may involve tolerance and withdrawal effects. In correlation with these public rhetorical strategies, the DSM has long designated the usage of psychoactive substances for medical purposes as *a priori* nonpathological (American Psychiatric Association, 1980, p. 163), while in medicine more broadly we are witnessing an ongoing shift from a model of health structured around ‘cures’ toward one of long-term risk reduction through ‘drugs for life,’ which Joseph Dumit labels ‘mass health’ (Dumit, 2012).

A myriad of countervailing rhetorics circulate in the popular press, however. Those presenting the sharpest arguments against ADHD diagnosis and drug treatment are even more colourful in their usage of metaphor. While the mainstream metaphors suggest the remedying of an ‘internal’ dysfunction, critical voices tend to emphasize the ‘subduing’ effects of the drugs, as did Bradley and other early researchers. Again, they refer back to medical technologies, but more controversial and fear-inducing types: the same *Washington Post* article cited above, for the eyeglass metaphor associated with ‘most specialists,’ also cites a skeptical clinician claiming that “you get

better behavior, but it's using medicine as a chemical straitjacket" (Allen, 2001). The straitjacket is perhaps one of the most powerful symbols of psychiatry as an institution of authoritarian control, and so it should come as no surprise that it gets put to rhetorical use by those seeking to present psychostimulants as the modern face of the same repressive apparatus. This metaphor is found primarily in the North American press, and occasionally in the British, but there similar arguments are often framed by a different regional metaphor, with tabloid opinion columnists and pundits regularly referring to Ritalin and other psychostimulants with the same phrase used to describe highly-sedating antipsychotics like Thorazine, as the "chemical cosh" (Womack, 2006).¹ Where these metaphors are presented in the context of opinion pieces or short stories reporting a particular claim from a public figure, they tend to indicate an anti-pharmaceutical, anti-psychiatric slant overall, and typically a relatively unsophisticated presentation of any scientific content. Sarah Womack's article in the *Daily Telegraph* cites the arguments in the British House of Lords by a public 'spokesperson' for neuroscience whose arguments about the effects of media on 'normal' minds have already been examined in previous sections, Baroness Susan Greenfield; here she reaffirms her technological etiology, arguing that the disorders are in fact caused by "too much time spent on a computer screen," and so drug treatment is unwarranted. Greenfield herself does not employ the 'chemical cosh' phrase, but Womack brings that in alongside a spokesperson for the Citizens' Commission on Human Rights (without explaining that it is a Scientology-affiliated anti-psychiatry organization). Clinical professionals routinely express concern that claims about the constructedness of mental disorders may be tied to more radical anti-psychiatry sentiment, and articles like these in the tabloid press suggest their fears are not unwarranted.

¹ A cosh, in turn, is a heavy cudgel, of the type one might use to knock a person unconscious.

Examples like the *Washington Post* article cited above, however, indicate that even where strongly critical voices appear, the complexities of diagnosis may be dealt with in a more careful and measured fashion. In this piece we find both the major technological metaphors, and an individual narrative (where Internet use figures in a largely incidental though positive way) where a quite heterogeneous set of actors gets brought into play: a scenario in which the label started to seem somewhat indefinite and increasingly oppressive, involving ‘comorbidities’ of diagnosis and accompanying polypharmacy (Clonidine, Prozac, and Pamelar alongside the boy’s Ritalin); a range of credentialed medical professionals expressing a diversity of opinions about diagnosis and treatment; concerned parents and teachers doing the same; all amidst online networks of lay expertise, industry-funded activist groups like CHADD, and not least, the spiritual ‘light’ of the family’s Quaker tradition. The main subject of the article, Andrew Foster, and his family begin to question the judgment of his doctor, eventually rejecting his advice to add Risperdal to the drug regimen (an ‘off-label’ recommendation, of course).² He chooses to discontinue all medication on the advice of a new psychiatrist, and in a followup published years later, readers are told that he finally found solace and productive employment by learning to become a dog trainer, where his classmates “treated him respectfully and not as someone with a disorder” (Seiss, 2011). This narrative is rich with signification in terms of the interplay between individual ‘plasticity’ and the ‘flexibility’ demanded within a socio-economic *milieu* (Malabou, 2008; Rose & Abi-Rached, 2013). What Andrew found as an alternative to the medications whose side effects became debilitating can, it seems, be read in two very different ways: as accommodating himself to the docility demanded of him as a labourer,

² As mentioned above, this was the same drug whose ‘off-label’ marketing earned its manufacturer, Janssen (a subsidiary of Johnson & Johnson) several large penalties in the United States (Muskal, 2012). Since this *Post* story was published, it was also the subject of a ‘disease-mongering’ controversy when it was revealed that Dr. Joseph Biederman failed to report at least \$1.6 million in consulting fees received from Janssen from 2000 to 2007, while at the same time promoting the diagnosis of pediatric bipolar disorder, for which ‘mood-stabilizing’ drugs like Risperdal were recommended (G. Harris & Carey, 2008).

conveyed inadequately to him by the primary institutions of education and medicine, instead by training himself to train such docility into the ‘pets’ of others, or, conversely, as a lesson in the psychic value of coming to know oneself, alterity, and the world more deeply through an *entre-deux* of multispecies interaction (in the manner of Haraway, 2007). No doubt Andrew would prefer the latter framing, as do I, in terms of doing justice to the range of agencies implicated within his odyssey of diagnosis, treatment, and subject-formation. Nonetheless, his narrative is equally marked by the many institutionalized powers amongst which behaviourally ‘troubled’ adolescents must negotiate their positions.

Psychostimulants are also routinely characterized in the press by analogy with other illicit drugs, in phrases like ‘kiddie cocaine’ reported and appropriated from the slang of users, or in references to ‘speed’ and recreational amphetamine use. As the young children diagnosed with attention disorders reach adolescence and high school, concern in the press over their usage of stimulant drugs shifts toward these tropes, and narratives of pharmaceutical diversion, misuse, and addiction. Individuals prescribed these medications but who do not take them as prescribed—either not at all, or only intermittently—are reported to be routinely selling Ritalin or Dexedrine in the schoolyard to others without prescriptions, who then often crush up the pills and insufflate them for greater euphoric effect. With these narratives, focusing on children as young as 8 or 9 for what we might call ‘alarm value,’ and more credibly and extensively on high school and university students, we find a different pattern of coverage, shifting away from concern about medical drugs as instruments for social control and instead as illicit forces to *be* controlled. At times this runs alongside skepticism about the medicalization of behaviour more generally, but more frequently this is presented as a regulatory challenge for medicine and education, of ensuring that only those legitimately diagnosed with a disorder may access these compounds. The implication is also often

made that there are qualitative differences in how individuals experience the drugs, that those with the disorder are able to ‘magically calm down,’ while “for people who are not hyperactive or do not have trouble paying attention, taking Ritalin and Dexedrine has the opposite effect: an immediate, euphoric and addictive high” (Walton, 2001). Such articles also typically flag concerns that prescribed stimulants may constitute a ‘gateway’ to illicit drug use.

The press must grapple with a range of thorny scientific and social causal questions in this regard, all under deadline and with limited resources. What do these drugs do to the brain, and is this different from the action of similar drugs under conditions of non-medically-sanctioned use? Does their medical use cause a propensity toward addiction - or should any observed increase in addictions amongst treatment populations be attributed instead to the disorder itself?³ Some professional sources cited in the media suggest that substance abuse among populations with ADHD has a grounding in the biology of the disorder, stretching neuroscientific plausibility to claim that use of nicotine, cannabis, or cocaine is an unwise but effective way to “activate the brain neurochemicals the patient is looking for... prescribed ADHD medication, however, meets their needs” more safely and effectively (according to Dr. Umesh Jain, head of the ‘Canadian ADHD Resource Alliance’ and a habitual source for Canadian journalists: White, 2008). There were several stories from around the world which reported patients taken off their medications in adolescence when ADHD was defined as a childhood disorder, who then quickly ‘turned to’ illicit substances. One whose headline asks ‘Prescription for Life?’ suggests it is good news that “guys in their late-20s and 30s who realise that off their medication they just fell in a heap” are then able to obtain a

³ This same problem of causality is even sharper in the context of anti-depressant or anti-psychotic medication and suicide, where the most serious of potential treatment side effects can also be readily explained as a symptom of the underlying disorder. David Healy has written extensively on this, suggesting that this rhetorical strategy bolsters the more fundamental statistical techniques by which the significance of adverse drug effects is occluded (Healy, 2012)

diagnosis of adult ADHD and continue psychostimulant treatment indefinitely (Christopher, 2007; cf. Dumit, 2012). Others suggest that ADHD “is not a valid diagnosis,” and that “huge numbers” of those admitted to addiction rehabilitation centres are there “because they have been given amphetamines as children” (Pryer, 2004).

There are no easy answers to be distilled by the consumer of this coverage, though on occasion individual articles or spokespeople may offer tempting candidates. How does the popular media interpret stimulant drug use? It ultimately depends where you look and who’s doing the using. Are we to believe that they remedy some underlying deficiency in the brain, and that this also explains why those with the disorder are drawn to structurally similar illicit drugs? Or are we to carry the analogy further, across the lines of demarcation between normal and pathological, licit and illicit, and see all stimulants as having the same effects on everyone? Being on speed helps you stay awake and be productive, as mathematicians, pilots, and long-haul truckers alike have understood since the days of Benzedrine.⁴ But is being on medication for ADHD the *same thing* as being on illegal ‘speed,’ be it from a clandestine laboratory or diverted from the pharmacy? These questions blend seamlessly with debates about the ethics of using these compounds for ‘cognitive enhancement,’ to which I return at the end of this section. When a college student insufflates some crushed Adderall and stays up all night writing a paper, there is persistent uncertainty about how this behavior should be named and framed. Is it illicit? Recreational? Treatment or enhancement? And how does having a prescription (or not) change matters, when amenable doctors are as easy to find online as the DSM checklists they use for diagnosis? In professional and in popular discourse, answers to these questions are diverse, and settled conclusions elusive.

⁴ Norbert Wiener is one of the mathematicians, along with Paul Erdos, whom we know often boosted their productivity with amphetamines during their mid-century heyday; there are almost certainly a many others from this era whose use of drugs for ‘neuroenhancement’ is not a matter of public record.

Gender presents further layers of complex reality and construction enfolded within this discourse. There really are more boys diagnosed with ADHD as children than girls, and this imbalance persists into adulthood. The full range of positions was expressed in my sample regarding the interpretation of this fact, however. Occasionally I found instances where this was dismissed as a matter of ongoing discovery - that girls were afflicted with the disorder at comparable rates, but remained undiagnosed because their symptoms were less overtly hyperactive. More often, the discrepancy was simply noted, without much speculation on its nature or meaning. Regularly, however, this fact was recruited for a masculinist strain of the same conservative commentary outlined above, with claims that ADHD signaled not only how “society was making a malady of boyhood itself,” but further that the educational system was structurally favourable toward girls, likely because so many teachers were female (C. Abraham, 2010). On this view, any failure to adapt was not on the part of boys but of institutions, and sources often went further in their implication that society has come to privilege excessively ‘feminist’ matters of concern: “‘What if we were drugging girls at the same rate?’ asks Jon Bradley, education professor at McGill University. ‘What if [the majority] of these prescriptions were being written for girls? ... There’d be a march’” (*ibid.*). Many similar articles and opinion pieces used the diagnostic imbalance in this fashion to link the critique of ADHD and psychiatry with a broader right-wing polemic against the supposed harming of boys by an ascendant, misguided form of feminism (Sommers, 2001).

As other researchers have noted, while fathers may be lingering on the sidelines of some such polemics, most of those who proclaim the irresponsibility of parents as a primary cause of disordered behaviour are tacitly or openly engaged in ‘mother-blaming’ (Singh, 2004). The *Daily Mail* was unsurprisingly a repeat offender in this regard, with one columnist, after first decrying the negligence of a mother in a ‘Glasgow housing estate,’ casting equal aspersions on her middle-class

counterparts: “Having children late in life, they are used to highpowered careers in offices that run like clockwork. They live in showhomes, a sticky finger has never sullied their high-gloss kitchen units. In this world of perfect presentation, shoes and handbags always match” (McAlpine, 2006). This brand of moralising about motherhood runs closely alongside complaints about the prevalence of single parenting, and the use of electronic devices as a means of distracting unruly children. All this typically adds up to a conclusion that ADHD, if it has any reality at all, is more typically a ‘cover’ for bad parenting, bad children, and an epochal shift away from personal responsibility. This reaffirmation of a traditional moral conception of childrearing has a certain contradictory quality about it: publicly blaming mothers for seeking psychiatric diagnoses, supposedly as a shield from blame.

This ignores both scholarly research and more nuanced personal narratives in the popular press, which suggest that public sentiment against ‘drugging children’ with stimulants remains high, and that parents of both genders often opt for these treatments in a context of powerful self-blame, guilt, and institutional pressure. Conversely, a couple of articles even touched on an unusual strain of debate about the validity of ADHD diagnosis in adult women, with one pediatrician quoted as being unsure whether “the Ritalin is treating their ADHD or allowing them to act as superwomen in our demanding culture” (Elias, 2005). Men may be diagnosed with the disorder at higher rates, but women are also confronted with a particularly challenging and worrisome discourse on ADHD, not only as the explicit focus of many narratives, but also as implied primary caregivers in many others. Around all of these public debates about diagnosis, scientific research and individual scientists get recruited in a range of ways, but the conclusions ultimately remain underdetermined and unsettled. One consequence is that short pieces presenting seemingly definite answers - whether based on science or skeptical of its conclusions - are some of the least credible. The best ‘scientific

journalism' in this instance is often less directly based on specific research, but rather longform narratives which aim to do justice to the complexities of individual positions. Andrew Foster's story discussed above is one good case in point.

This should become clearer as I proceed to examine some of the most frequently-cited public 'authorities' on the specific question of media effects. Actor-network theory has always proposed that as some agents put themselves in a position to be seen and interacted with as 'spokespeople' for another set of actors, translating their heterogeneous concerns into a common 'interest' (Callon's *interressement*), they become 'obligatory passage points' (Callon, 1986); this model was extended to account for situations in which the passage points are less properly *obligatory* and more plural, with differing passage points allying different but oft-overlapping sets of actors around shared boundary objects (Star & Griesemer, 1989). The situation I examine here is more of the latter type than what one would find in 'settled' or less-publicized scientific controversies, with a diverse set of actors from different social realms seeking to make their own work into an "obligatory point of passage for the whole network of participants" (*ibid.*, p.390). This occurs on the level of the media outlets themselves, as they compete with one another for audience share and advertising dollars, the basic feedback cycle which drives a tendency toward controversy-mongering in general, and the proliferation of skeptical stories about ADHD and drug treatment specifically. The popular media in aggregate may at the same time be considered the primary passage point for public knowledge about science and medicine, as well as the uncertain entities around which they cohere; be they considered boundary objects, epistemic things, or matters of fact in-the-making, it is through the process of mediation in both the Latourian and vernacular senses that they become matters of public concern. Journalists are crucial to this process of translation, of course, but in developing their stories they often have recourse to another fine exemplar of the passage point, the shared *list of*

sources. By some combination of self-publicizing, professional public relations, and institutional outreach efforts (often by research universities), certain figures come to be regional or international spokespeople on a topic, regularly contacted as a source by journalists working for different media firms. In concluding this section, I describe two such sources for the matter of ADHD and media, one from within the institutional bounds of ‘science proper,’ pediatrician and epidemiologist Dmitri Christakis, and one more popular figure, Edward Hallowell.

Christakis is the lead author of the best-known research paper examining the relationship between media use and attentional problems (Christakis et al., 2004), at present cited by over 500 other articles, and widely reported in the popular media. This paper employed data from the United States National Longitudinal Survey of Youth, produced starting in 1979,⁵ to argue that high levels of early television exposure before the age of 3 were associated with higher levels of hyperactivity and attentional problems later in life. In many ways the coverage of this study is a classic tale of journalistic hyperbole. While the study does clearly argue for a correlation between these two variables, it is explicit that it cannot be interpreted causally without further evidence—though the authors argue that by focusing longitudinally on early television exposure between ages 1-3, before diagnoses of ADHD are typically made, this limitation is ‘mitigated.’ Their survey measure of ‘attentional problems’ also means they “have not in fact studied or found an association between television viewing and clinically diagnosed ADHD” (Christakis et al., 2004, p. 711). The measured tone of the published paper stands in sharp contrast to the headlines, which almost invariably reported this as scientific evidence that children who watched more television were more likely to

⁵ NLSY79 is the data set giving information on parents, and NLSY-Child gives linked information on children of its female respondents. Christakis et al. draw a sample of 1278 children who were between 6 years 9 months and 8 years 9 months of age during the NLSY-Child ‘survey waves’ in 1996, 1998, or 2000, and analyse the correlations between ratings on a ‘hyperactivity subscale’ of the survey’s ‘Behavior Problems Index,’ and reports of their TV viewing while one and three years of age, as well as NLSY79 measures of ‘maternal self-esteem and depression.’ (Christakis et al., 2004, p. 709).

develop ADHD. Though the study was based on an analysis of survey research collected by the US government, there was little discussion in the press of the methods, measures, or data sets involved. Instead, there were regular injections of ‘neuro-essentialism’ even into coverage of this relatively mundane epidemiological analysis, with the ubiquitous circuit-metaphor for the brain called into service to frame this study as evidence that television was ‘rewiring the brains’ of developing children (“Turn off the television,” 2004, “TV ‘rewires’ developing brains, researchers fear,” 2004). What was acknowledged in the paper as a limitation of its data set was spun into a strength, moreover, as the researchers’ admitted lack of any data on the content of programming watched was put in service of an argument that content - paraphrasing McLuhan - was merely the ‘juicy meat’ which distracted the ‘watchdog of the mind’ while the structure of the medium exerted its influence.

“Unrealistically fast-paced visual images typical of most TV programming may alter normal brain development,” journalists reported, and “it doesn’t matter if toddlers watch cartoons or educational shows, the rapid-fire stimulus of the medium itself sets up mental habits that may feed an inability to focus and concentrate later on” (*ibid.*). In a majority of media narratives where they occurred, ascriptions of a causal role to ‘screen time’ were enrolled within a broader critique of the medicalization of attention disorders, with suggestions that diagnosis and drug treatment were misguided responses to a set of behaviours better treated by reducing media consumption. This connection was proposed by opinion columnists, by ostensibly scientific ‘popularizing’ spokespeople like the aforementioned Susan Greenfield, and by writers who used it as a foil to the mainstream consensus (*ie.*, ADHD is a ‘real disorder,’ specifically not caused by ‘too much TV’). In Christakis’ research and in coverage thereof, however, we see that this connection is by no means a necessary one, and that in fact arguments about a potential environmental factor like the media in shaping ADHD need not imply skepticism that it is a real, brain-based disorder. The paper itself cites the

shortcomings of an ‘emphasis on structural or neurologic features of the central nervous system,’ and invokes neuroplasticity to argue that ‘the types of visual and auditory experiences’ in early childhood “may have profound influences on brain development” (Christakis et al., 2004, p. 708). Again, the methodology of the paper itself is survey-based and in no way neuroscientific, but for contemporary discourses claiming to speak scientifically of behaviour, the brain is the most obligatory passage point of all. While the birth and flourishing of the disorder may have been in an environment where on the whole neither children, parents, nor society were to blame (A. Lakoff, 2000), and while commentators continue to conceive its medicalization along those lines, Christakis’ study and its popular mediation suggests how brain and blame may now be more readily coupled together.

Scientific research offers only limited and intermittent evidentiary contributions to public debates on the epidemiology and etiology of mental pathologies; spokespeople with varying sorts of credentials deploy these to ally the interests of diverse actors (individual and collective) through a particular interpretation of the causal relations between heredity, biology, environment, society, responsibility, and any other agencies posited as effective. A skeptical, socially-conservative strain of opinion privileges personal and parental responsibility above all, while a similarly skeptical but more progressive variant privileges environmental and social factors. Christakis presents another alternative, whereby neuroplasticity and the media as environmental factor work to bind together responsibility with a mainstream neuro-essentialism. His claims in the professional literature about the mixture of hereditary and developmental factors in this disorder, as in many domains of human biology and pathology, are measured and plausible. In its wider mediation, however, this research buttresses the notion that even if ADHD is reducible to a fault within the brain, and not a matter of

personal moral culpability, individuals are nevertheless called upon in contemporary discourse to ‘manage their brains’ and avoid potential negative influences for themselves and their children.

Edward Hallowell, described in his promotional materials as ‘the first to name adult ADD,’ is another widely cited figure on the disorder in general, but particularly on the relationship between the media and its rise to supposed ‘epidemic’ status. In interviews for the print and broadcast media, and in a popular book (first published in 1995 and recently re-issued: E. M. Hallowell & Ratey, 2011), he has been the most prominent advocate for the idea that this disorder was a ‘metaphor for modern life,’ and a maladaptive response to our ‘hypermediated world.’ While identifying personally as having ADD, he argues that there are a range of possible responses to a high-stimulation environment. Those who have what he calls “attention-surplus disorder” thrive in this context, readily multitasking and meeting the demands of a ‘knowledge-based economy,’ while others cannot cope and manifest symptoms of inattention or hyperactivity requiring treatment (Owens, 2006). In this sense, they present a more complex and speculative account of the role played by biology and environment, even offering a third potential response in another piece about ‘the lure of data:’ ‘pseudo-attention deficit disorder,’ whose sufferers

do not have actual ADD, but, influenced by technology and the pace of modern life, have developed shorter attention spans. They become frustrated with long-term projects, thrive on the stress of constant fixes of information, and physically crave the bursts of stimulation from checking e-mail or voice mail or answering the phone (Richtel, 2003).

Thus they again argue for a responsiveness of the brain to its developmental *milieu*, but in a fashion which potentially encompasses far more than even the growing ranks of those diagnosable with ‘actual’ clinical ADD, to also include those who (like Christakis’ survey populations in fact) spend a lot of time with technology and find sustained ‘deep’ attention difficult. They also describe their theory in metaphorical terms drawn from neuroscience, and make connections to drugs of abuse, in

this case referring to the media themselves using a familiar trope, with Hallowell's co-author claiming "it's like a dopamine squirt to be connected... it takes the same pathway as our drugs of abuse and pleasure" (*ibid.*). A self-described 'radical moderate,' across Hallowell's media appearances he is careful to avoid any radical skepticism toward the disorder, or broader anti-psychiatric critiques, while at the same time proposing a sociotechnical explanation for the disorder's apparently increasing prevalence by relating it to proliferating information technologies, and to associated changes in the pace and pattern of human interactions.

This brand of 'moderation' is an eminently marketable one, insisting on the one hand that there is an organic basis in the brain for both 'real' attention deficit disorders and responses to media, but on the other that there is equally a proliferation of 'pseudo-ADD' cases, in which somewhat maladaptive responses on the boundary between normal and pathological are presented in association with an excess of mediated 'connectedness.' Once again this involves a call to manage the brain, and to govern interactions with technology prudently. This particular call is directed for maximum appeal, however, toward three distinct audience communities simultaneously. First, as an avowed but evidently high-functioning sufferer of 'true' ADD, Hallowell affirms the core beliefs of those who accept the reality of the diagnostic label (as well as the usefulness of pharmaceutical treatment), and presents his own narrative testimony of achieving a productive life in spite of this disability. He conversely offers anecdotal analysis of cases at the far opposite end of the spectrum, those with 'attention surplus disorder' who flourish in a world of ubiquitous information technologies, achieving high levels of productivity and satisfaction. And for the broad swath of those who routinely find themselves distracted, inattentive, or otherwise matching the diagnostic criteria for ADHD, and also feel that technology is in some sense responsible, *without* believing themselves afflicted with a true mental disorder, Hallowell likewise offers solace - with a range of

‘self-help’ books, motivational speeches, and workshops. He trades on his credentials as a professional psychiatrist and former Harvard Medical School faculty, much like the other ‘expert’ sources cited in my sample, and his brand of ecumenical popular mental health discourse recruits scientific and medical claims in a variety of ways. Yet it also goes far beyond them, achieving prominence not simply by public relations campaigns, but by supplementing psychiatric categories with an interpretive schema tailored to appeal to the widest potential audiences, whereby all of the psychological dispositions with respect to technology may ultimately be turned into ‘assets.’ Psychological, social, and economic well-being are tied in with coming to understand oneself through one of these dispositions or temperaments. The loop thus encompasses not only those diagnosed with a pathology, but those who are in fact *supra*-normal, and a wide spectrum in between.

This appealing message even attracts a second-order reflexivity on at least one occasion in the press, as one piece from 1997 which profiles Hallowell at the second Attention Deficit Disorder Association conference goes on to ponder the looping public construction of the disorder:

ADD may be just a story they tell themselves, a comforting label for behavior they have been unable to explain or control... When we feel that something undefinable, unnameable is wrong with us, there is great relief, great validation, in simply naming our condition. And there is even greater relief in the discovery that others share it. (Glusker, 1997)

This fascinating long piece, another from the *Washington Post*, bemoans a world where “an entire subculture has grown up around ADD,” a world of ‘well-meaning professionals’ alongside ‘charlatans and profiteers,’ marketers, conferences, and online forums, in which “science and medicine have mingled so thoroughly with capitalism and 12-step revivalism that it’s hard to tell where the boundaries are anymore” (*ibid.*). These blurred boundaries between capitalism and science, medicine and communities of lay expertise, necessarily imply a blurring of the boundaries between normal and pathological. The question becomes not so much whether ADD is ‘real’ or ‘just a story,’ as when and in what contexts the labels have which effects. What does the proliferation of different

‘gurus,’ each pushing their own take on the disorder and its rising prevalence, imply for the reality of the condition? What are we to make of the different ways that people come to understand themselves by accepting or resisting a diagnostic label? What of the pharmaceutical marketers who rebrand an appetite suppressant (Obetrol) as Adderall, turning a variant of the same old Benzedrine into a new patented ‘blockbuster’ drug?⁶ These are only some of the concerns that circulate through this piece and many other longer articles, and they are ultimately better descriptions for their refusal to offer easy answers. Broadly, the media coverage does show that “totalizing critiques of scientific reductionism” (Pickersgill, 2009) are misguided. Biological reductionism, neuro-essentialism, and a ‘techno-somatic’ ethos (*ibid.*) privileging neurotransmitters and genes over social epidemiology are all in evidence throughout the discourse, but while these stances may be appropriately considered ‘mainstream’ within psychiatry, they circulate in the press alongside many other heterogeneous ways of understanding behaviour. There are firm defenses of the former views, but even among the professional experts and authorities cited, there are pockets of resistance and uneasy mixtures of social construction with organic causation.

The last major issue that becomes apparent in relation to Hallowell’s work is neuroenhancement. While regarding many cases of attentional problems as ‘pseudo’-disorders, arguing for a substantial role played by media in influencing these patterns of behaviour, and admitting that there is no firm diagnostic line by which to distinguish the truly pathological from its ersatz form, Hallowell remains an advocate for the usefulness of drug treatment.⁷ This seems to be a

⁶ Roger Griggs, the Shire pharmaceutical executive who introduced Adderall, coined the new name after purchasing the original manufacturer of Obetrol, reportedly after spending some time “fiddling with snappy suffixes:” “All. For A.D.D. A.D.D. for All. Adderall. ‘It was meant to be kind of an inclusive thing,’ Mr. Griggs recalled.” Indicating some of the tensions negotiated by such executives, however, Griggs also states that “he strongly opposes marketing stimulants to the general public because of their dangers. He calls them ‘nuclear bombs,’ warranted only under extreme circumstances and when carefully overseen by a physician.” (Schwarz, 2013b).

⁷ A passage from Hallowell’s website gives a fascinating explication of this position, again through technological metaphor, while also underlining the distinct economic orientation of his approach: “I tell the child [with ADHD] that

relatively common position amongst more moderate clinicians: that while a lack of diagnostic reliability and potential overdiagnosis are concerns, pharmaceuticals remain a safe and effective treatment for disorders of attention. Yet framing matters thusly raises serious questions about who might benefit, and whether the uses of the drugs should in fact be limited to the treatment of a specific disorder. As prescriptions for psychostimulants have steadily increased, so too has their usage by those without prescriptions. This recalls many of the same concerns from an earlier ‘amphetamine epidemic’ in mid-twentieth-century America (Rasmussen, 2008), and I have already detailed some of the worries presented in the media about the diversion and potential abuse of these drugs for ‘recreational’ purposes. Yet as I have already suggested, many contexts of use confuse the boundaries between licit and illicit, as well as recreational and self-enhancing aims. Should everyone who feels beset with Hallowell’s brand of ‘pseudo-ADD’ have access to psychostimulants?

Particularly toward the end of my sampling period, a number of lengthy pieces in major newspapers considered precisely these uncertainties. One from the *Observer* opens with the experiences of a college student at Harvard, who obtained an Adderall prescription by describing symptoms “that he knew were typical of the disorder” (but without apparently believing himself truly afflicted), and used it throughout his postsecondary education to work long hours while maintaining an active social life (Talbot, 2009). Like most other stories on this topic, it focuses heavily on university students, and cites research indicating that from 5% to upwards of 35% of them may have engaged in ‘non-medical’ use of prescription stimulants to enhance cognitive

he is lucky in that he has a race car for a brain, a Ferrari engine. I tell him he has the potential to grow into a champion. I tell him (assuming it is a he, but he could just as easily be a she) that with effort he can achieve greatness in his life, and then I tell him about the billionaires, CEO’s, Pulitzer Prize winners and professional athletes with ADHD I’ve treated over the years. But I also tell him he does face one major problem. While he has a race car for a brain, he has bicycle brakes. I tell him I am a brake specialist, and one of the many tools I can use to strengthen his brakes is medication. I remind him he will have to do much more than take the medication to strengthen his brakes, but, if we’re lucky, the medication will help him in that effort.” This critical essay was written in response to a New York Times piece which he accuses of ‘sensationalizing’ the dangers of Ritalin (E. Hallowell, 2012).

performance (B. P. White, Becker-Blease, & Grace-Bishop, 2006). It raises some concerns about addiction and potential other side effects from non-prescribed use, while referencing a new strand of the same plaintive conservatism as discussed above. One commentator bemoans the young people of today working with “their laptop, their iPhone, and their Adderall,” placing stimulant drugs in company with other behaviour-reshaping technologies seen to signal the end of authentic human experience. Of course, as already discussed in earlier sections, similar complaints with respect to new media are perennial, but this phrasing signals how drugs constitute media in themselves. This article is far from a moral-panic piece about the irresponsibility of college students, however. It assembles together a diverse collection of individuals using a variety of drugs – the former including university professors as well as a professional poker player, the latter including not only traditional stimulants like Ritalin and Adderall but newer ‘wakefulness-promoting agents’ like modafinil and unregulated ‘ampakines’ like piracetam. The professors, for their part, are cited not just for their ‘performance-enhancing’ usage of these drugs, but also for having published bioethical arguments for the legalization and acceptance of such uses (Sahakian & Morein-Zamir, 2007).

Clinicians and researchers in the article confront the dilemma of the ‘worried well:’ those who may not be truly stricken with a disorder (though they certainly know the symptom checklist and may match some entries on it), but who feel they are not performing optimally in school or at work. Should they be prescribed stimulants? There is research suggesting at least short-term enhancement of cognitive function regardless of whether one has an authentic disorder, and little evidence for how to distinguish an authentic pathology from normal challenges. Should universities permit their students to use such drugs, or should they be considered ‘cheating?’ These are thorny bioethical questions, cutting right to the heart of what it means to have an authentic core of subjectivity apart from its myriad technological modulations. Following on earlier debates about the

advent of what Peter Kramer called ‘cosmetic psychopharmacology,’ aimed toward becoming ‘better than well’ (P. D. Kramer, 1997), one of the ethicists cited in this article concluded in another *Nature* piece (‘Professor’s Little Helper’) that the best response to concerns about neuroenhancement is legalization and physician oversight. Doctors could supervise the usage of stimulants and other ‘cosmetic’ psychopharmaceuticals, just as cosmetic surgeons provide their services, without necessarily requiring a diagnosis of pathology. Demand is simply too high, and continued prohibition is becoming another front in a doomed War on Drugs: “it would be difficult to stop the spread in use of cognitive enhancers given a global market in pharmaceuticals with increasingly easy online access. The drive for self-enhancement of cognition is likely to be as strong if not stronger than in the realms of ‘enhancement’ of beauty and sexual function” (Sahakian & Morein-Zamir, 2007). What emerges from all of these discussions is that the demands of a Western knowledge economy are perhaps some of the most potent forces in the construction of attention deficit disorder. The behaviours thus labeled are problematized in the context of failure to meet those demands, first in the educational system and then ultimately in the context of flexible, entrepreneurial, precarious post-Fordist labour. And productive interaction with a range of media technologies is defining feature of not only labour but leisure in this era.

What was initially regarded as a transient childhood manifestation of psychic conflict and as a defect of ‘moral control’ is now perceived globally, but not universally, as an internal disorder of the brain; drugs thought to subdue behaviour by repressing a flood of unconscious impulses are now more likely to be seen as treating a dopamine deficiency. Diagnosis is so uncertain and deeply social, however, that a multiplicity of powerful popular etiologies circulate alongside more authoritative professional claims. Hallowell, Christakis, and many others give voice to a popular sentiment that diagnosed cases of the disorder are only the most overtly maladaptive responses to

the accelerating pace of our technological society. These popular passage points use the press as a platform for *interestement* to ally scientific claims (and claims of scientificity) with a commonly voiced suspicion that “we’re all a little ADD” in the age of the Internet. This idea takes shape in a variety of distinct forms, and the matters of concern vary. The consequences of blurring the line between normal and pathological with respect to ADD may be to regard stimulants as a dangerous ‘chemical straitjacket’ being applied to ordinary misbehaviour, as a valuable treatment occasionally misused and overprescribed, or even as a universally valuable resource for those trying to get ahead in school or in business.

A broadening of the potential market for psychostimulants evidently finds allies in the pharmaceutical industry. Even an article which is generally open to the usage of these drugs for personal enhancement reports with some skepticism about the cycles by which pharmaceutical companies seem to construct disorders alongside the drugs that treat them, and how the same compounds may be defined and marketed in changing ways, over time and simultaneously as well. Like many amphetamines, the patented formulation of Adderall was first marketed in 1960 as Obetrol for what is now considered a ‘side effect,’ its suppression of appetite. Modafinil is described in the press in terms of a similar phenomenon, “mission creep,” whereby a drug first marketed to narcoleptics is within a few years also approved for a broadening category of individuals with ‘excessive daytime sleepiness’ unrelated to narcolepsy, and eventually ‘shift work sleep disorder’ (Talbot, 2009).⁸ The last disorder seems to establish the outer bounds of the limit between normal and pathological—a long-acting new stimulant is accepted for use in keeping at bay a quite normal bodily response, not only for fighter pilots and emergency personnel but for any night-shift

⁸ Similar concerns have been examined extensively in scholarly work on the pharmaceutical industry (Dumit, 2012; Healy, 2002; Shorter, 2008).

labourers with difficulties remaining awake. Beyond ‘shift work sleep disorder,’ the only remaining untapped market for these drugs is for pure enhancement, without any pretense of a disorder.

Whether such uses will ultimately be judged as acceptable may have less to do with bioethics or government regulation than by market forces. As one analyst puts it,

if you’re a company that’s got 47 offices worldwide, and all of a sudden your Singapore office is using cognitive enablers, and you’re saying to Congress: ‘I’m moving all my financial operations to Singapore and Taiwan, because it’s legal to use those there’, you bet that Congress is going to say: ‘Well, OK.’ It will be a moot question then. (Talbot, 2009)

Debates about ethics and morality may eventually start to seem somewhat beside the point. Claims like this herald a future in which potential side effects and addiction are to be weighed not against the risks of leaving a disorder untreated, but instead against those of falling behind economically—be it as an individual, a corporation, or a nation.

The potential negative effects of new media on attention spans may one day be understood in much the same way as those of an unhealthy diet: with stimulants prescribed like statins, as chemical alternatives to a difficult regimen of behavioural self-policing for the management of risk. In this way, such drugs promise a mode of self-fashioning which Joseph Dumit calls ‘better living through chemistry’ (Dumit, 2012). As a *Wall Street Journal* piece arguing in favour of legal neuroenhancement contends, one advantage of that framing may be to keep insurance companies off the hook. Let the diagnoses remain for those with ‘genuine’ illnesses that deserve to be covered, while at the same time allowing the worried well to undergo a “neurochemical nose job” if they’re willing to foot the bill: “we should be tough about limiting the insurance burden for such drugs to those who do have serious illness. But if people pay for safe psychopharmacology, that should be their choice” (S. Satel, 2003). A world of chemical enhancement for normal populations would be a world where insurers - whether private or public – cease to be a meaningful check on the expansion

of drug sales. The line between sick and well would no longer determine who had access to psychostimulant drugs, only who would be paying for them.

Presently, however, one needs a prescription. Across my sample I found many questions were being raised over whether all of those being prescribed such drugs really had disorders in need of treatment. One final figure who has recently become a passage point for media discourse on ADHD is *New York Times* reporter Alan Schwarz. After spending a few years covering the NFL's head injury crisis, he turned to ADHD and psychostimulant drugs as another perceived mental health crisis where medical ethics came into question. In a series of lengthy pieces in the *Times*, Schwarz provided a platform for a number of skeptical voices and narratives concerned with the pharmaceutical treatment of ADHD. One piece opened with a quotation from Keith Conners, creator of many widely used diagnostic instruments for the disorder, who questioned the increasing numbers of those diagnosed, stating that "The numbers make it look like an epidemic. Well, it's not. It's preposterous... This is a concoction to justify the giving out of medication at unprecedented and unjustifiable levels" (Schwarz, 2013b). The article went on to offer critiques of the marketing strategies for branded drugs like Ritalin and Vyvanse as having "stretched the image of classic A.D.H.D. to include relatively normal behavior like carelessness and impatience, and has often overstated the pills' benefits" (*ibid.*). Another piece by Schwarz chronicled the struggles of a college student named Richard Fee, who began using Adderall diverted from friends' prescriptions to fuel last-minute study sessions, then acquired prescriptions of his own, escalating his dosages to the point that he developed psychotic symptoms, and ultimately committed suicide (Schwarz, 2013a).

Yet others raised questions about the prescribing of stimulant drugs to toddlers (Schwarz, 2014), or to low-income students who may not have an authentic disorder. One doctor he cites as having diagnosed many children with A.D.H.D. called the disorder "'made up,' and 'an excuse' to

prescribe the pills to treat what he considers the children's true ill – poor academic performance in inadequate schools” (Schwarz, 2012). Schwarz's writings add up to a powerful critique of a new amphetamine epidemic, closely entwined with an enormous expansion of A.D.H.D. diagnoses, with the implication that despite some valid cases of the disorder, in a great majority of others it serves to paper over more serious issues. Where pressure to perform in educational or business settings is great and social support or individual aptitude are insufficient, psychostimulants present a risky solution. He presents a world in which it is difficult to draw lines with any certainty between treatment, neuroenhancement, and reckless, compulsive drug use. Is he performing a valuable service by driving public concern about these issues, or unfairly stoking resentment against the medical establishment and contributing to the stigmatization of those diagnosed with the disorder? I incline toward the former view, but there have been multiple open letters and online petitions suggesting the latter (T. Brown, 2013). In many ways it is a matter of perspective.

It should be recognized, however, that concerns about overdiagnosis and overprescription do not contradict the reality that many of those diagnosed and treated with stimulant drugs have found their cognitive performance, subjective life satisfaction, and self-conception all greatly improved. Despite the dichotomous framings which often reign in popular media, we would do better to adopt a ‘dappled’ view of causation in psychiatric disorders (Kendler, 2009), accepting of heterogeneity in the factors which give rise to mental illness, in the subjective experiences of their symptoms, and in the pathways of treatment. Disorders seem likely to be caused not exclusively by internal brain dynamics or external causes, not by inborn, inherited temperament or by ‘rewiring’ through psychological development and technology, and instead by all of the above in varying mixtures. This is perhaps the most general normative conclusion to be drawn with regard to a discourse on ADHD and the media wherein a great many themes and images circulate without any

overriding general character. The topic occasions considerable debate regarding social construction, neuroplasticity, biological essence and personal responsibility; these debates might proceed in a more reasoned fashion informed by a dappled view of causation, ‘both/and’ rather than ‘either/or.’

Instead, there is one pattern of claims which is most often presented in the press as that of a psychiatric ‘mainstream’: that ADHD is a genetic, heritable disorder, caused by some uncertain but distinctive brain dysfunction, and precisely *not* by ‘too much TV’ or other sociotechnical factors. There are also claims presented which hold quite the opposite, that the disorder *is* in fact a product of our social context, with accelerating change and technological ubiquity playing a significant role in its genesis. Such notions may be woven together by opinion columnists in a surly conservative rebuttal of the ‘brain blame narrative,’ with excess media consumption and ADHD figured instead as intertwined effects of lax parenting and declining ‘traditional values.’⁹ Yet despite stories which often present them as such, brain-based accounts need not be mutually exclusive with those which take media effects into account. We are naturally inclined toward the artificial, toward becoming-cyborg, with neuroplasticity positioning us as subjects open to modifying our brains and minds through interaction with technology (Haraway, 1990; Clark, 2003). This idea, that technology may be changing us, for better or for worse, has been central throughout this project and its many digressions. I believe it is, as do many others whose work I have discussed. It is not a determining influence, however, but only one of many within a matrix of heterogeneous social and biological factors which play a role in shaping cognition. While necessary, this dappled view of causation in mental illness and mental health is far from sufficient in ensuring reasonableness. The brain may

⁹ Drug treatment may be figured likewise along more progressive lines as an abdication of social responsibility, often as a spreading but particularly American propensity to treat collective problems as individual disorders: far easier to medicate supposedly pathological brains than remedy the myriad deficiencies of the public school system, or formal and informal support systems for the poor, or poverty itself.

readily be worked into dubious opinion columns, allied with the same old mother-blaming and TV-blaming; Susan Greenfield, again, has made an art of leveraging neuroplasticity for fearmongering purposes.

Sociotechnical concerns may also be united in more palatable but still somewhat muddled ways with brain-based accounts, as I found in the writings of public spokespeople like Christakis and Hallowell. Borrowing a phrase from Catherine Malabou, both are exemplary of the “neuronal form of social and political functioning,” the subtext for much of the discourse in my sample, and a form which “deeply coincides with the current face of capitalism” (Malabou, 2008, p. 10). In addition to his own studies, Christakis publicizes psychological research on ‘critical periods’ not as an alternative to but as complementary with standard psychiatric accounts of ADHD, suggesting that we ought to avoid allowing too much screen time to children in order to preserve their attention spans but casting no doubt on the reality of the clinical diagnosis. In this sense our brains and the brains of our children are conceived as objects of careful management and governance. Likewise with Hallowell, who renders ADHD as a phenomenon which may in certain contexts take on the appearance of pathology, but in others may be harnessed to achieve the supra-normal levels of functioning demanded in a post-Fordist economy of ‘flexible accumulation.’ He too raises no skeptical questions about the disorder itself or drug treatment, even as he suggests we may all be a ‘little bit ADD’ today. Being that way might be the best way to get ahead in an era which demands constant shifting of focus and readiness to adapt to change. If you should find it hard to cope and your doctor should prescribe you some Adderall, well, that’s just another tool to ‘strengthen your brain.’ As many of the voices in my sample suggest, even skepticism about the disorder or about the firmness of a division between normal and pathological states of attention need not preclude the deployment of psychostimulants as technologies of behavior. Whether ADHD in general or any

specific case is believed to be 'real,' the improvements brought about by these drugs in at least some aspects of cognitive functioning are difficult to dispute. Whether they are worth the accompanying risks, or whether the gains are truly lasting and meaningful, remains quite uncertain and shaped by individual contexts of use.

Hence we arrive by a multiplicity of looping paths at a set of conclusions mostly quite amenable to our present 'spirit of capitalism.' Our public discourse circles around the twin labels of ADHD and of the drugs we use to treat them, offering antagonistic approximations vying for authoritative status, but no consensus as to their real nature. Whether caused by the biology of the brain, by new media, by social change or some mixture of these and other factors, disordered attention is seen as an epidemic in our time. Even those who decry it as manufactured still label it as such. As I have contended throughout this project, with further research from both the sciences and humanities, we may yet arrive at a fuller understanding of the causal interconnections between minds and media. For now we are scarcely beginning to understand the world of mediated cognition. In my sample, however, the only scaling back of consumption advised with some regularity is of television viewing and other forms of screen time for children. When it comes to the demands of education, production, and the technologies thereof, whether they may cause pathologies or they stand innocent of the charges, the solution is uniformly presented as adaptation.

One may adapt by cultivating one's facility for hyper-attention, or one's niches for deep attention and reflection; by understanding oneself through a diagnosis of ADD or by resisting one; one may cultivate mindfulness and time-management techniques or employ psychostimulant drugs. One person may adopt all of these techniques or none to adapt and cope. To renounce technology, however, is becoming less and less an option. Working with computers and the Internet can hardly be refused on the grounds that it makes you feel scattered and shallow, at least not without sharply

limiting one's options. What is to come in a near future when, as Satya Nadella, current CEO of Microsoft, has put it, the "true scarce commodity" will be "human attention" (Egan, 2016)?¹⁰ Will the devices to which we pay attention leave us vapid, or will we develop new forms of hyperintelligence alongside hyper-attention through our interactions with machines? Will the market for psychostimulant drugs stay limited to those defined as sick, or open up to the potentially limitless group of those seeking better adaptation and faster advancement within this economy of scarce attention and flexible accumulation? Whatever the coming years may hold, however uncertain these answers may be, attention deficit disorders are mediated phenomena on many levels, as with so many other aspects of our cognition, and they stand among the most paradigmatic pathologies of our time.

¹⁰ In an ideal closing vignette for this section, I distractedly clicked away to the *New York Times* site in the midst of writing, and after reading something entirely unrelated about blizzards just happened to come across this piece, only the latest of many op-eds in which the theme of an attention economy is conjoined with the assertion that "our devices are rewiring our brains" (Egan, 2016).

Epilogue

“Take the human, for instance. Of course, it is no longer a calculating entity which could easily be morphed into silicon chips. But it is certainly not a subjective, reflexive, intentional, embodied unity either. Not only has its cognition been distributed, situated, but it is now shared with many intellectual technologies to the point where studying a human is studying a field of forces and transfers of documents, instruments, ideographies, through a collective of similarly distributed fellows . . . The engineering dream was to morph the human into a rational machine. The humanist counterdream was to recover an intentional, reflexive and coherent carrier of values. The result is a rather bizarre cyborg that resembles neither the machine nor the human.”

— Bruno Latour (1996, p. 302)

The foregoing project can be read as an effort to comprehend this destabilized, decentered human subject, outlined by Latour and effected by contemporary technoscience. What is it to be human, to exist as a human consciousness in this era of information technology? Are we but one more class of complex information-processing systems? We confront a series of opposing conceptions, each potent but unsatisfying: first of all between the calculation of the Hobbesian dream, that cognitivism for which ratiocination is the hallmark of humanity, and nothing other than computation itself, ‘easily’ realized in silicon chips; and at the other pole, the mythic unity of a humanistic, intentional, value-driven subject. There are various associated points of antagonism as well: between *homo economicus* and bounded rationality; between representational planning and situated, embodied actions; between the ‘natural’ mind, bounded within the brain, and the extended, distributed mind, enacted through external media. Each perspective has some merit, and in other ways proves inadequate. Human existence today, at least in the developed world, is poised between these incompatible alternatives. We call upon on heterogeneous resources in fashioning our objective selves, in coming to know what we ‘really are.’ Some of these are scientific, some spiritual, some technological and some philosophical, some moral and some medical. Only rarely do these

come together in some fully consistent, coherent self-conception; more commonly, they are layered upon one another, giving rise to unanticipated and aporetic hybrids.

As specified from the outset, in closing this sequence of interconnected studies I offer few definite explanations or settled conclusions. Recapitulating the key themes, however, an overall coherence can be established, a constellation or 'galaxy' in the style of Innis and McLuhan. I have sought to explore the deep interconnectedness of human minds and technological media: across history and looking toward the future, in science and in popular imagination, in theory, fiction, practice, and reality. I opened by framing the problems of media studies and of science and technology studies as closely aligned and mutually complementary. Analysing McLuhan's theory of media in parallel with accounts of technology from actor-network theory and elsewhere in STS, I considered the dilemma of technological determinism. We can best understand media effects not as determining, isolinear causal forces, but as emergent from a matrix of sociotechnical interactions. Indeed, technology itself should be understood as a name for this social domain of human-machine interactions, rather than as a conceptual stand-in for the machines themselves. This opens the way to thinking of media effects as something other than deterministic, external influences upon society. Such effects are not new and corrupting influences on natural human cognition, but on the contrary we are naturally predisposed to generate such artificial effects. Deep attention and logical reasoning are not transhistorical universals of advanced thought but products of a culture long founded upon the printed word. The risk, as Bernard Stiegler and others have noted, is that new media forms may be supplanting our capacities for sustained attention and reasoning, thereby undermining our ability to serve as enlightened political subjects. Hence media may ultimately undermine democracy: a long-

standing concern of communications studies through the twentieth century and its many propaganda campaigns.¹¹ While occasionally exaggerated, this worry is not unfounded.

Human rationality is always-already shaped by interaction with technology, and fearmongering about new media is often groundless. Nevertheless, valid concerns remain as to specific technologies and techniques which, whether by design or unplanned emergence, do have the potential to bring about collectively deleterious effects on our cognition. These may not be making us all shallow or stupid, but they are changing how we think and interact. Sometimes these effects can be positive: as when we feel more connected to a more diverse array of people, perhaps contributing more regularly to charitable causes, or when our behaviours are ‘nudged’ in prosocial ways by carefully-constructed media campaigns exploiting what we know about heuristic processing. In other instances these can be more troubling, as when we find correlations between social media use and anxiety disorders, when researchers on social media sites experiment with the information they present to see if it can make us depressed, or when other forms of ‘A/B’ testing are aimed toward simply capturing more of our attention and disposable income. Science is an indispensable resource in understanding these kinds of effects, but it is by no means exempt from them in its own practices, nor is it free from complicity in efforts to deliberately create them. In attempting to shed some light on these dimensions of the science-media interaction, I first considered the role of technological metaphor in theorizing the mind, particularly those metaphors drawn from telegraphy and subsequent related ‘control technologies.’

Metaphor is in some sense foundational to cognition itself (Hofstadter, 1979; Leary, 1990; Ricoeur, 1993), and so it should come as no surprise that it plays a reflexive role in our

¹¹ In fact, if we look at Plato’s *Republic* once again and its conception of the ‘noble lie,’ the democracy-undermining capacities of mediated narratives are already recognized then, and actually rendered as a positive; this line of thinking was taken up in the twentieth century as well, most notably by Leo Strauss and Walter Lippmann.

understanding of our own minds as well. This is so both in popular discourse and within the specialized vocabularies of the cognitive sciences. Metaphors drawn from the latest and greatest technologies have long held appeal in characterizing the mind. First came clockwork automata, which Hobbes so appreciated and Babbage later tried to concretely implement as an analogue to cognition; then, more importantly for my purposes, we began to draw upon telegraphy. This technology, the first which allowed messages to be transmitted faster than a messenger could travel, and which first allowed us to translate information into electricity, came to furnish some of our most potent generative metaphors for understanding the mind. Along with the co-constitution of modern science and technology in the telegraph—and its associated ‘musculature,’ the railways—came the arrival of neuroscience, and for the first time we began to conceive the substrate of the human psyche by analogy with a network of electrical wires. Yet the role of the telegraph was not simply to furnish metaphors. The analogy points toward a more fundamental homology, in that both brain and telegraph are signaling systems which employ electric current as their medium of transmission. Beyond generative metaphors, the telegraph system provided a set of concrete techniques for conceptualizing and quantifying information and communication. It opened the way to understanding and manipulating these phenomena as immaterial constructs abstracted from the particular media in which they were instantiated—a necessary condition for the development of cognitive science. This dual valence of technology, as source of both tropes and tools which together drove the sciences of mind, only increased with the development of the telephone, and then the transition from mind-as-telegraph to mind-as-computer.

The information theory of Shannon, Weaver, and Wiener represents the clearest line of filiation connecting all three domains, diffusing as it did from the technical journals of the Bell System to the status of invaluable resource across diverse scientific disciplines. The history of

cognitive science is inseparable from the history of control technologies, from the era of cybernetics through to our own. Systems like the telephone, analog calculating machines, and then later general-purpose digital computers, constituted ‘things to think with’ in myriad senses: at once rhetorical and practical tools, extensions and models of cognition; systems for conducting research, but also attracting funding, and eventually for disseminating it as well. This mode of technoscience was always closely allied with military and corporate interests, and in that sense the deployment of control technologies as both practical tools and models for understanding the mind is inseparable from ‘closed-world discourse’ (Edwards, 1997) and the administration of a ‘control society’ (Deleuze, 1992). The human subject understood as information-processor is equally a subject susceptible to modeling, monitoring, and new modes of governance, a subject becoming-data in multiple ways, for multiple purposes. At the same time, however, key actors within cybernetics and later cognitive science imagined different futures, while some actively worked against the subordination of their efforts to the aims of military and industry. Norbert Wiener’s career was paradigmatic in this regard, tracing a path from the development of cybernetics in wartime work on weapons-control systems, to later anti-nuclear activism and concern for the effects of automation in the labour market. As in other areas, scientists affiliated with cybernetics and cognitive science espoused diverse positions regarding such social issues. At no time, however, could they be considered to have been conducting ‘pure’ science wholly divorced from these concerns.

Turning to the history of cognitive science, I presented a case for the deeply mediated nature of the field, and its founding upon the mind-as-computer analogy. Its practitioners also sought to move beyond analogy, however, toward concrete, manipulable models. These would make scientific discourse on the mind respectable again: devices and programs could be built to act in ways traditionally ascribed to the mind, but with their inner workings far more accessible to investigation,

and evidently lacking any mystical or transcendent elements. Metaphors drawn from computing technology are the clearest common thread binding together the disparate movements grouped together under the name of cognitive science. Yet as the kinds of machines employed by researchers in these fields changed, so too did their fundamental conceptions of mind. As models were built which began to demonstrate behaviours reasonably characterized as intelligent, the metaphor started being taken increasingly literally by some cognitive scientists and by the popular press. This meant that the research was both highly influential, and subject to a boom-and-bust cycle of mediated 'hype' and subsequent backlash, when it became apparent that these systems had significant shortcomings as compared with human cognition.

The first such cycle came with the cybernetics movement, which began as a series of conferences on 'circular causal systems,' and eventually became a focus of popular fascination. First came the fire-control systems like Wiener and Bigelow's predictor, but then later 'clicking brains' like Ashby's Homeostat and mobile robots like Grey Walter's tortoises were subjects of hyperbolic journalism, suggesting that the capacities of such machines might soon equal or surpass our own. These were analog technologies, processing continuously variable elements with a combination of electronic circuits, vacuum tubes, and mechanical actuators. They operated not in isolation, but as part of a total reconceptualization of mind, behavior, and society, responding directly to the political context of the Cold War and aiming to simultaneously cultivate global peace and psychiatric health. It was founded upon a new metaphysics, seeking to understand both biological and mechanical systems through a unified vocabulary of control and information flow. The core of this cybernetic movement developed through the interdisciplinary Macy Conferences, and in successive 'waves.' Its technological apogee came with Frank Rosenblatt's Perceptron, an early connectionist model which sought to capture quite abstract and general properties of cognition, discriminating between stimuli

through parallel processing with neuron-like elements. While a small contingent of researchers would go on to pursue a ‘third wave’ of cybernetics, cognitive science properly speaking would form in reaction against the Perceptron, with Minsky and Papert’s attack on its capabilities constituting a touchstone for this new era. The mainstream of scientific research into the mind shifted from pursuing specially-constructed analog systems as models of cognition, to developing software models on general-purpose digital computers. Public interest followed, with this next generation of ‘thinking machines’ and their advocates also becoming subjects of enthusiastic press coverage. The notion of a cyborg, originally a tangential offshoot, eventually came to eclipse the scientific movement altogether in the public consciousness; for most people, the term ‘cybernetics’ is now more likely to conjure up an image of the Terminator than of Norbert Wiener.

With the rise of cognitivism came the notion of multiple realizability. This meant that its models were seen as something more than copies. Rather, cognition as such was understood as an abstract phenomenon—a phenomenon of information, for that concept remained just as essential—which could be implemented either through evolution, in biological systems, or by design, in artificial systems. This view was most clearly espoused by two giants in the nascent fields of computer science, artificial intelligence, and cognitive science, Herbert Simon and Allen Newell. They proposed the physical symbol system hypothesis: that the capacity for autonomously manipulating and generating strings of representational symbols according to definite procedures is necessary and sufficient for a system to be considered ‘thinking.’ Hence despite its rather rudimentary capacities, they considered their General Problem Solver an authentically intelligent software program, and made dramatic predictions for the future of their field. More important to the ongoing success and development of cognitive science, however, was the ‘heterogeneous engineering’ of actors like J.C.R. Licklider, who focused ARPA spending on this area, as one prong

in a strategy for promoting computer science, developing networking architectures for interconnecting digital computers, as well as sharing and collaboratively processing scientific data in new ways. The dreams of producing authentic artificial intelligence by the 1970s did not come to fruition, of course. What was produced instead was a different conceptualization of the mind, as an abstract information-processing system, alongside different modes of interacting with computers, and with other humans by way of computers. Gone were the days of laboriously punching paper, then waiting for one's ration of time to feed it through and see if it produced any meaningful output; now codes in higher-level languages could be entered on screens and output viewed in something approaching real time. The experience of working with such systems gave rise to what has been labelled 'cognitivism,' the orthodoxy in AI and cognitive science for which Newell and Simon were the strongest advocates. Computers became tools for not only comprehending but fully instantiating the one underlying, universal process which produces cognition.

As the predictions for artificial intelligence failed to materialize, ARPA became DARPA, pruning its research budgets and directing them toward projects with direct military applications. The field entered a proverbial 'winter,' and the pursuit of thinking machines ceased to command the interest of researchers, funding agencies, and mass media as it once had. Research developed instead along two distinct lines. One focused on 'expert systems,' and other techniques for employing computer software as extensions of cognition, to offload some of the information-processing burden in tasks involving both a higher-level reasoning component and a brute-force search component. The other focused once again on developing models of cognition, reformulating psychological inquiry in terms of information-processing, and giving rise to cognitive science as such. Many of these models were developed according to heuristic programming approaches akin to those of 'good old-fashioned AI;' this era also marked the return, however, of connectionist

approaches to cognitive modelling, which aimed for greater neurobiological plausibility, mirroring the microstructure of cognition in the design of their programs. These two approaches are often seen as standing in opposition to one another. As I tried to emphasize, however, many of the most interesting models, beginning with Hofstadter's Copycat, incorporated elements of both. More recent models like Blue Brain seek to incorporate as much neuroscientific detail as possible—though some contend still not enough—in producing a system that carries out cognition by simulating a portion of a biological brain on supercomputing hardware. Others, like SyNAPSE, strive to build new hardware architectures based on stripped-down neuron-like elements, with public controversies ensuing as to the appropriate level of detail required to deem a system an adequate model of the brain.

Massively powerful supercomputing-based projects like the Blue Brain may seem far removed from electromechanical gadgets like Ashby's Homeostat, and they are indeed exponentially more complex. At the same time, however, they continue the same entanglement of tools and theories that has held sway in the sciences of cognition since the time of cybernetics. Their departure point is the generative metaphor of mind-as-computation, but they seek to instantiate this in tangible and testable systems, pursuing a unified conception of biological and artificial cognition as information-processing. Current systems like Blue Brain and SyNAPSE return more closely to the biomimetic origins of this pursuit, incorporating analog elements as well. Other projects continue to pursue a more abstract and representational form of modelling, like Douglas Lenat's long-running Cyc project; still others like SPAUN lie somewhere in between, leveraging distributed computing and open-source production. Any modeling of the mind now also benefits from dramatic improvements in neuroimaging technologies over the past decades, though questions remain about the validity of that data, hence also about the prospects for success of projects based upon it. Across

the years and the heterogeneity in its techniques, what unifies cognitive science is the vision that cognition can best be understood through technological media. At the same time, it enacts highly mediated new forms of computer-assisted cognition itself, rendering mind once again a scientifically respectable concept by reformulating it in ways that would be impossible without the external supports of digital hardware and interactively-coded software. Cognitive science also helped cultivate new modes of interaction with and by networked computer, which would eventually diffuse out across the scientific community and later come to transform mass culture. It can thus certainly stand as a resource in understanding the effects of media upon the mind, but only if we understand its origins in reconceptualizing the very notion of mind by analogy with computational media.

Following my consideration of this history, I turned to the study of media effects, tracing the history of that discourse from wartime propaganda research to McLuhan and Gerbner's correspondence on the neurobiology of media. In this I uncovered the groundwork for a conceptualization of these effects which saw them as neither powerful deterministic forces, nor insignificant limited ones, but as nevertheless playing a subtle, crucial role in shaping our ways of thinking and acting. From McLuhan I take the notion that formal properties of media exert influences on a subconscious, neurological level, and from Gerbner that the wider patterns of media content cultivate psychological dispositions in often-unrecognized ways. While in principle we are far better equipped today to scientifically investigate both sorts of cognitive effects, in practice I found that most research into these questions was of the traditional psychological type, with relatively little work being done on these topics using brain imaging or computational modelling. One could certainly be forgiven for thinking this was not the case, given that contemporary discourse abounds with neuro-buzzwords in describing these long-standing concerns.

Commentators speak about aggression and media violence in terms of amygdala activation, or compulsive media use as a matter of seeking out ‘dopamine squirts.’ There are indeed some brain imaging studies, as I discussed, which do lend some support to these speculative connections. The evidence is ambiguous and susceptible to multiple interpretations, however. One small imaging study, for instance, which found similarities between patterns of brain activation in video game players and in those who have been administered psychostimulant drugs, has been cited in support of entirely opposing views: some say it shows how games can *cause* attentional problems, others how it can *treat* them.

In the end, neuroscience and modeling support the building of a consensus founded primarily on traditional psychological studies. There is such a consensus building for moderate but significant effects in terms of media violence promoting aggressive cognition, as well as media use magnifying problematic traits and behaviours in those predisposed to addiction, anxiety or depression. As commentators are often prone to do, however, in taking the step further and claiming that violent media produces violent acts, or addictive media produces mental pathology in otherwise normal individuals, we move beyond what is supported by scientific evidence. The more reasoned view, I contend, is to understand these effects through the prism of cultivation analysis and bounded rationality. Given our finite cognitive resources, it is only too understandable that we draw as much, if not more often from mass media than direct personal experiences in forming beliefs and making decisions. Thereby extending the sphere of our cognition may have positive or deleterious effects, and a good deal of further research from all branches of the cognitive sciences will be needed to better comprehend the processes underlying both sorts of effects. In another sense though, digital media is already everywhere within the cognitive sciences, transforming each stage of its research processes. Presenting a phenomenon to a subject immobilized within a brain scanner is

almost invariably presenting them some form of electronically mediated content, while consumer virtual-reality technologies and techniques adapted from gaming open up new possibilities for imaging research. At the same time though this makes it difficult to distinguish the specific effects of media within the brain—what is one to compare them against?

The cognitive sciences are now pervaded by the Internet and, to a lesser extent but perhaps more surprisingly, electronic gaming. Services like Amazon's Mechanical Turk have taken on roles in furnishing research subjects, leading to some issues with quality of responses but also some advantages in moving beyond the traditional populations of psychology undergraduates employed in many studies. Other sites and games used as vehicles for social interaction have used A/B testing not only to find more effective ways of presenting content to drive user base growth and retention, but to produce specific emotional states in their users. The nascent integration of academic psychological research with this brand of experimentation has been fraught with controversy (Kramer et al., 2014; Hughes, 2014). New ethical norms will need to be worked out when it comes to manipulating massive populations of perhaps-unwitting users. Even if it is disclosed somewhere in a lengthy End-user License Agreement that users may be subjects of research, does clicking 'I agree' truly constitute informed consent? Still, the possibilities for electronically gathering data on much larger and more diverse samples cannot be ignored. Along with the large user bases of online games, and the usage of game-like environments within brain scanners, the field of 'brain training' also sees digital games as potentially therapeutic technologies. On the one hand, addiction research confronts electronic gaming as a possible source of compulsive, problematic behavior; on the other, brain training seeks to leverage our tendency to play these games for hours on end, using them as vehicles for complex mental tasks which may forestall cognitive decline, particularly in older adults. Debates over brain training also signal how the Internet has transformed the dissemination and

critique of psychological research. Many of the games claiming cognitive benefits can be downloaded commercially, while a few others have sought formal approval as medical devices. Cognitive scientists have spoken out against such products, particularly the former group, with many signing one online 'open letter' contending any generalized benefits for these games would be minor, and that evidence for their value is very limited. Another group with comparable credentials, however, responded with its own open letter, arguing that despite some marketing hyperbole, the approach has merit, and statements to the contrary may have a chilling effect on funding and development for new games that might have greater benefits. Both petitions continue to gain new signatures online from the communities of interested researchers.

Such online spaces are crucial to the making of scientific consensus today. They render this making all the more visible and public, with matters that would once have been confined to closed meetings and conferences now aired for global audiences. The debates which proceed in blogs and online forums are in a sense even more public than the contributions in formal peer-reviewed venues, given that many of the latter are only available at an unreasonable expense for those without access through an academic institution. At the same time, turning to such venues may lead to accusations of scientific 'impropriety.' In the domain of mental health where I turned in my final section, the Internet has played an enormous role in sustaining debates regarding the status, definition, and particularly *re*definition of mental disorders. No revision of the American Psychiatric Association's Diagnostic and Statistical Manual has been subject to more public scrutiny than the most recent fifth edition, producing an array of open letters and petitions online. The APA even instituted an official online public comment period. Communities of patients rose up in vocal dissent against proposed changes, for instance with the elimination of Asperger's syndrome as a diagnostic category separate from autism spectrum disorders. Changes in psychiatric definitions of

mental disorders have always produced changes in the self-conception of patients, and more recently in the financial realities of their insurance coverage. Diagnosis has never operated in a vacuum, and the clinic has never been isolated from society. But now clinicians, researchers, and policymakers cannot avoid these realities or the mass-mediated public debates connected to them. When people are classified and reclassified, they react. We cannot avoid confronting this reality, as new modes of communication have granted worldwide influence to discursive communities forming around these reactions, whether in support of or against mainstream medical opinion.

In presenting my research on attention-deficit disorders, I sought first to emphasize the changing face of psychiatric diagnosis, and the heterogeneous medical conceptualizations of such conditions over the years. The Whiggish histories popular within the professional literature present a vision of linear progress from a preliminary conception in terms of defective moral control with George Still, through the byways of psychoanalysis to an evidence-based biomedical conception of a disorder grounded in the brain and treatable with stimulant drugs. Against this view, I emphasized the significant differences across the multiplicity of labels leading up to ADHD, and the kinds of hybrid understandings which defied easy categorization—for instance that of Charles Bradley, pioneer of amphetamine therapy, who understood these drugs' effects in psychodynamic terms. On my reading, alongside the circuitous history of our labels for the disorder, the career of the labels attached to such drugs is of equal importance. How marketing departments brand, governments regulate, and the public responds to drugs like Ritalin and Adderall, all play an enormous role in the social constitution of attention disorders. Past conceptions of both disorder and drug often figure implicitly into contemporary discourse, at times contrasting and at other times blending the moral and the medical, the psychosocial and the neurobiological. In my final chapter I turned to one more

looping cultural pathway, considering the question of whether media may be causing attention disorders, as taken up in the mass print media.

In the popular press today we find a good deal of moral-panic discourse about a generation possessed of ‘so much media, and so little attention span,’ spending its days ‘fixated by screens and nothing else.’ This project began with my idly noticing such articles on a fairly routine basis, and when it came to systematically sampling the literature, I found a great deal more of it. I hypothesized that there would be considerable discussion of the potential causal role played by media in rising diagnosis rates for attention disorders, particularly in children, and at the same time that there would be majority support for brain-based conceptions of the disorder, as well as for pharmaceutical treatment. Essentially all these hypotheses were confirmed, with 35 percent of all articles I sampled giving voice to the idea that computers or TV may be causing attention disorders. But there was more heterogeneity of opinion than I expected, and I found that support for the view that media could cause ADD was often aligned with anti-psychiatric, anti-pharmaceutical sentiment. This in turn was coupled with a traditionalist conception of parental and individual responsibility. Particularly in the tabloid press, I found a vision of society in which neglectful parents leave their children in front of televisions or computers for hours a day, leaving them incapable of paying sustained attention to anything else, then use the disorder as alibi for their failure, and pharmaceuticals as a ‘chemical cosh’ to suppress it. Only a small percentage of articles explicitly rejected the idea that technology could be a causal factor, typically using this vision as an anecdotal foil to mainstream psychiatric opinion. Such articles rejected ‘too much TV’ or bad parenting as spurious folk epidemiology, and insisted that attention deficit disorders were the result of inheritable brain abnormalities. Many pieces I found slotted neatly into this bifurcation, supporting either a biogenic psychiatric mainstream, or a conservative-minded total rejection of same. Yet I also found a

greater than expected number of multifaceted, interactional accounts, with other articles implying—rightly, in my view—that the causes of the disorder are multiple, blending neurobiological factors with patterns of human-machine interaction as well as other possible social and environmental influences.

Attention deficit disorders are now an eminently *public* construction, with a range of ‘gurus’ and would-be experts seeking to position themselves as obligatory passage points for journalistic discourse about ADD/ADHD. I highlighted a few of these whose work was particularly relevant to the question of technology as potential cause. Dimitri Christakis’ 2004 study is one of the most often-cited points of evidence for the theory that excess television exposure may cause problems with attention span, and he is a widely cited authority on the topic who has more recently investigated the same theory with animal models (both papers discussed above: Christakis *et al.*, 2004; Christakis *et al.*, 2011). He has furnished some of the best evidence in apparent support for the view that ‘screen time’ should be carefully managed among young children. The American Academy of Pediatrics suggests that it be avoided altogether prior to the age of 2, yet Christakis has somewhat surprisingly criticized the rigidity of this guideline and suggested that interactive screens may be less harmful than watching television. Edward Hallowell was another whose writing I discussed, a more popularizing, guru-type figure who simultaneously contends that ADD/ADHD is a real biomedical phenomenon meriting pharmaceutical treatment, but at the same time that the disorder is a ‘metaphor for modern life,’ and we are all experiencing ‘pseudo-ADD’ as a result of the increasing pace of technological interaction in our lives. There can no longer be any question of renouncing all these attention-sapping technologies for most of us, but rather we must adapt. The solutions are manifold: we may keep our children away from screens, we may ration our time and shut off our devices for set periods, we may cultivate deep attention by reading printed books or we may

accustom ourselves to rapid-fire hyper-attention with digital games. We may eventually even come to see psychostimulants as bearing the same relation to digital media as statins for cholesterol, a pharmaceutical alternative to a challenging regime of behavioural self-policing. When the Internet has you distracted but you can't afford to disconnect, just pop an Adderall!

The biomedical conception of mental health is strengthening its consensus, but remains deeply contested in our time, and it enters into hybrid forms with earlier conceptions of human nature. One may come to accept that it is the brain and not the psyche or soul which is afflicted in a disorder like ADD, while still holding that neglectful mothers and social liberalism are to blame. One may believe that the disorder is fictitious, but that at the same time applying the diagnosis to struggling children and prescribing them stimulant drugs is the best way to help them when the public school system cannot. Even in its 'purest' form, understanding behavior as a matter of neurobiology does not imply fatalism about our somatic destiny. Instead, the brain/mind becomes an object of governance and self-management in new ways. Attention is a resource to be managed and paid out efficiently, its continuing supply to be maintained by guarding ourselves and our children from negative media effects, amongst other possible influences on the plastic brain. Much the same can be said of memory and executive function in older adults, except in that case digital media, in the form of brain training games, are instead often presented as the solution rather than the problem. An understanding of neuroplasticity over the course of human development means that neuroscience can provide little guidance on its own for understanding behavior. Instead it adds one more layer of mediation to our objective selves, another circuit through which the manifold influences of society, culture, and technology must pass in the production of our subjectivity. Hence whether normal or pathological, the mind is best understood in terms of a 'dappled,' multifaceted causal structure (Kendler, 2012). Effects do not flow in one direction from the organic brain

outward to psychology and society, nor vice versa, but endlessly circulate back and forth in a strange loop.

Human cognition, then, is a phenomenon mediated in manifold ways. It is by nature poised to expand and extend itself through artifice. Technology does not affect us as some dangerous novel force which has corrupted our natural state, but rather as the precondition for our distinctive mental capacities. As the external media available to extend our capacities change in character, so too do our ways of thinking, interacting, and reflexively understanding ourselves. New media should not be greeted with fear and panic. But older media have shaped our minds and cultures in a diversity of profound, now-naturalized ways—above all the habits of deep attention and linear argumentation which accompanied print. We must attend, therefore, to what is gained and lost as new technologies drive change. They shape our habits of thought and action, our bounded rationalities, in a great many ways: both bidden and unbidden, for better and for worse, in designed ways and accidental ones. While in many ways mistaken, there is certainly some merit to the view that digital technologies lend themselves to a broad-ranging but shallow style of cognition and interaction (Carr, 2011). The Internet makes it easier than ever to gather together enormous quantities of data and different sources of opinion, but more difficult to spend hours delving deeply into a single book. It has never been easier to entertain instantaneous two-way communication with friends and family from around the globe, but it is perhaps harder than ever to have a long, profound, uninterrupted face-to-face conversation between two people. Such effects operate in an aggregate, statistical, and far from deterministic fashion. But at the same time, we must not let a fear of technological determinism blind us to their reality. As McLuhan proposed, there is never any inevitability when it comes to the media, provided that we are willing to contemplate what they are doing. The effects of

media are social, contextual, and can be actively resisted, but not if we imagine they are mere tools subordinated to human intentionality.

A similar story can be told about the part played by science in forming popular conceptions of mind and self. It plays a potent but by no means determinative role in our objective-self-fashioning, with talk of neurons and computational metaphors recruited alongside a disparate collection of other conceptual resources in understanding what our minds ‘really are.’ As we have seen, within scientific discourse such questions are far from settled, but nonetheless popular narratives clutch to any fragments of fact bearing upon these timeless matters of concern. We do so, moreover, in a biased fashion. Recursively reinforced by patterns of both scientific publication and popular media coverage, we are more likely to worry about a study which suggests new media are ‘rewiring our brains’—perhaps stimulating our amygdalas and causing aggression or sending ‘dopamine squirts’ to make us addicted—than we are to be reassured by a null result. Even when we recognize that these findings are only provisional and the conclusions underdetermined, neuro-buzzwords lend them an outsized importance, while a precautionary principle implies that we should nonetheless govern our conduct and that of our children to guard against potentially negative media effects. At the same time, media are reshaping science in profound ways. Computers have always played a crucial role in the theories and practices of the cognitive sciences, functioning by turns as communications technologies, suggestive metaphors, and concrete models. They pioneered a mode of human-machine interaction which has become commonplace. Now, the Internet is coming full circle and altering research practices on multiple levels: how data is collected, with systems like Mechanical Turk, as well as social media and online games; how researchers communicate with one another and with the public; and how in turn the public responds to and interacts with the community of scientists. Accounts of the public understanding of science founded upon one-way

diffusion of information have never been particularly satisfying, but digital media are rendering them patently absurd.

As I have suggested from the outset, science and media have long been closely intertwined, hence the problems of media studies and of science and technology studies are often one and the same. These interdisciplinary fields have approached them on different terms, however, and throughout I have tried to bring these discourses together in a constructive fashion. From media studies, I borrow the conviction that communications technologies act upon us and upon our society, shaping our interactions through both the forms they take, and the patterns of content they present. In their contemporary digital forms, these technologies cannot be produced without the efforts of scientists, and no effort to understand their effects can ignore the potential contributions of cognitive science. McLuhan, particularly in his later work, recognized this, but incorporated cognitive science in a rather rudimentary and uncritical fashion with his account of hemispheric biases caused by certain media forms. From science and technology studies, then, I draw the necessary antidote: close attention to the processes by which scientific knowledge is produced. This historically-informed understanding of the cognitive sciences completes the loop and allows us to recognize them not as neutral sources of factual information, but as intimately connected with the ascendancy of digital media. Does the fact that scientists have long understood the mind in terms of computation undermine their ability to investigate the effects of computers and other new media on the mind? I do not believe that it does. Rather, with that context in mind, we can better gauge the strength of different scientific contributions to these debates. We can also recognize the difference between research actually informed by neuroscience and computational modelling, as opposed to traditional forms of psychological inquiry which have just been garnished for popular consumption with buzzwords from these domains. With the 'critical neuroscience' movement, I call for a

judicious skepticism in interpreting new scientific findings about the mind and brain, treading a middle path between equally uncritical brands of acceptance and rejection.

Perhaps one day we will have a mature neuroscience which can map and model the human brain, supporting a cognitive psychology which could then demonstrate the fundamental homology between mind and computation. Rather than thinking in terms of human-computer interaction, we could fulfil the cybernetic dream and develop a science of interactions between differing species of information-processing systems. This day is far off, and may never come. But neither can we hold that it is impossible in principle. There is nothing ineffable or transcendental about the individual human mind nor the social collective which might definitively rule out such a comprehensive accounting. Just as neuroscience does not reveal our somatic destiny, though, such an understanding would not eliminate but enrich our concepts of mind and society. Even as information-processing systems we are embedded in multiple cultures and multiple economies. Our internal essence, if there is such a thing at all, is ontologically secondary to our interactions with others, human and nonhuman. Media technologies exert their effects not by some unique causal power, but by patterning these interactions in different ways, and by adding new nonhuman actors into the mix. These are often greeted with panic, particularly focused on children and adolescents. When we consider the complexities of both new media technologies and scientific research into their effects, we find a rather more varied picture than the one presented in narratives of moral decay and mental decline. Along with some evidence for positive effects, we also confront the end of an era when either manipulation or edification were 'broadcast' to the masses by elite opinion leaders. Instead we live in a time of many-to-many bidirectional communication, presenting new possibilities for diversity, debate, and resistance, but also for implementing a distributed society of control. The cognitive science of media has as much to do with developing techniques for producing targeted

effects—whether for business or government, as the McLuhanesque approach did in the 1960s—as it does with answering basic questions about the mind and its artificial extensions.

Foremost among the economies in which we find ourselves embedded today, I have argued, is the attention economy. The primary actor in this economy today, for better and for worse, is Google. We are scarcely beginning to understand the ramifications of the change it has wrought. As James Gleick puts it, “the merchandise of the information economy is not information; it is attention. ...Attention is what we, the users, give to Google, and our attention is what Google sells—concentrated, focused, and crystallized” (2011). Google’s users are at once consumers, products, and integral components within its technology. Every time you conduct a search or create a link online, you are fueling its artificial intelligence. “Search and advertising thus become the matched edges of a sharp sword,” with the perfect search engine reading your mind and producing the answer you want, while the perfect advertising engine reads your mind and feeds you a link to purchase the product you need. Advertising thereby becomes ‘virtuous’ (*ibid.*). We only see products we actually want, and advertisers only pay for their ads when users actually pay attention to them.

One does witness this on Google from time to time: the top search result is also the top advertising result, and of course the company has always assured us that the neutrality of the former algorithm would never be compromised by the latter. Yet this is the exception rather than the rule, and as Gleick goes on to contend, “if our interests and the advertisers’ were perfectly aligned, they would not need to pay. There is no information utopia. Google users are parties to a complex transaction, and ... we are not always witting parties” (*ibid.*). While we must not greet new media with panic, our eyes must be open to the ways our attention is being nudged, governed, monetized, and commoditized through them. Facebook’s study on emotional contagion is a signal in this regard, with the underlying method of A/B testing just one of the ways that corporate actors can seek to

generate targeted effects through digital media. The strategy can be very innocuous or it can be very troubling. It seems like a pleasant idea when targeted at improving the quality of interaction in online communities, but less so when it aims to collect more of our money in ‘microtransactions’ or keep us glued to a particular site for unreasonable lengths of time. It seems like a different and altogether more worrisome thing entirely when it aims to manipulate political outcomes. Repressive regimes already enlist armies of paid commenters to lend the appearance of public consent in online forums; manipulating social network feeds to skew perception of the opinions held by even our authentic friends and family is certainly a possibility on the horizon.

Democratic societies must be alert to these troubling developments, but it is not as if older print and broadcast media were without their own distinctive practices for manipulating public opinion and attention. Ultimately, new computational media are extensions of our cognition much like the others which have shaped our minds and societies across history. What is most distinctive about them is how they integrate and hybridize all prior media forms, encoding them in a universal language of information, and offering new possibilities for not just extending but delegating portions of our cognitive load onto autonomous nonhuman agents. Their effects are diverse and enormous in potential scope. I have tried over the course of this study to suggest some directions for further research, and ways of improving our scientific understanding of these effects: looking beyond the childhood-harm paradigm and media violence or addiction, toward a conception of cognition as always-already embedded in and distributed across systems of symbolic media. The rise of these new media is unlikely to spell the end of Enlightenment values or liberal democracy, nor will it lead us to produce or become hyperintelligent, disembodied consciousnesses—certainly not anytime soon, at least. Technological media are one factor among many influencing the genesis of human subjectivity, a crucial but by no means overriding element within the spectrum of causal forces making us who

we are. We cannot let the spectre of determinism lead us to ignore their effects. Equally, we should not let the shortcomings of a computational account of mind lead us into a wholesale rejection of the discourse, insisting as Robert Epstein recently has that “your brain does not process information, and it is not a computer” (R. Epstein, 2016). We certainly have not reached the point of a literal, quantifiable homology between human brains and digital computers. It remains at root a generative metaphor, yet it is no *mere* metaphor. Cognitive science has labored for half a century to make it more than that, and we cannot in good faith allege that information-processing theory has produced “few, if any, insights along the way” (*ibid.*), just because we can’t find discrete bytes of binary code circulating through our neurons. If, one day, a new system of metaphors supplants that of computation and information-processing in our understanding of the mind, it will likely do so by building upon, rather than demolishing, the edifice of computational cognitive science. In the meantime, digital computers are on their way to becoming as much a part of our cybernetic cognitive assemblages as brain tissue and natural language themselves.

Bibliography

- Abraham, C. (2010, October 19). 2 million prescriptions for attention deficit disorder written for kids, a jump of 43%. About 75% of them for boys. Are we medicating a disorder or treating boyhood itself as a disease? *The Globe and Mail*, p. A1. Toronto, Can.
- Abraham, T. H. (2012). Transcending disciplines: Scientific styles in studies of the brain in mid-twentieth century America. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, 43, 552–568.
- Adams, F., & Aizawa, K. (2008). *The bounds of cognition*. Malden, MA: Blackwell Pub.
- Adee, S. (2009, November 23). Cat Fight Brews Over Cat Brain. Retrieved October 7, 2014, from <http://spectrum.ieee.org/tech-talk/semiconductors/devices/blue-brain-project-leader-angry-about-cat-brain>
- Agre, P. (1993). The Symbolic Worldview: Reply to Vera and Simon. *Cognitive Science*, 17, 61–69.
- Agre, P. (1997). *Computation and Human Experience*. Cambridge: Cambridge University Press.
- Agre, P., & Chapman, D. (1990). What are plans for? *Robotics and Autonomous Systems*, 6, 17–34.
- Alac, M. (2008). Working with Brain Scans. *Social Studies of Science*, 38, 483–508.
- Alač, M., & Hutchins, E. (2004). I See What You Are Saying: Action as Cognition in fMRI Brain Mapping Practice. *Journal of Cognition and Culture*, 4, 629–661.
- Albergotti, R. (2014, July 1). Furor Erupts Over Facebook's Experiment on Users. *Wall Street Journal*. Retrieved from <http://www.wsj.com/articles/furor-erupts-over-facebook-experiment-on-users-1404085840>
- Albus, J. S., Bekey, G. A., Holland, J. H., Kanwisher, N. G., Krichmar, J. L., Mishkin, M., ... Tononi, G. (2007). A proposal for a decade of the mind initiative. *Science*, 317, 1321.
- Allen, A. (2001, March 18). The Trouble with ADHD. *The Washington Post*, p. W08.
- American Psychiatric Association. (1980). *Diagnostic and statistical manual of mental disorders (3rd ed.)*. Washington, D.C.: APA.
- American Society for Cybernetics. (n.d.). Summary: The Macy Conferences. Retrieved July 15, 2014, from <http://www.asc-cybernetics.org/foundations/history/MacySummary.htm#Part2>
- Anderson, C. A., & Bushman, B. J. (2001). Effects of Violent Video Games on Aggressive Behavior, Aggressive Cognition, Aggressive Affect, Physiological Arousal, and Prosocial Behavior: A Meta-Analytic Review of the Scientific Literature. *Psychological Science*, 12, 353–359.
- Anderson, C. A., Shibuya, A., Ihori, N., Swing, E. L., Bushman, B. J., Sakamoto, A., ... Saleem, M. (2010). Violent video game effects on aggression, empathy, and prosocial behavior in Eastern and Western countries: A meta-analytic review. *Psychological Bulletin*, 136, 151–173.
- Angell, M. (2004). *Drug money: how pharmaceutical companies deceive us and what we can do about it*. New York, NY: Random House.
- Anguera, J. A., Boccanfuso, J., Rintoul, J. L., Al-Hashimi, O., Faraji, F., Janowich, J., ... Gazzaley, A. (2013). Video game training enhances cognitive control in older adults. *Nature*, 501, 97–101.
- Anonymous. (2013). Camelot! Retrieved from <https://nucambiguous.wordpress.com/2013/11/30/camelot/>
- Anselme, P., & Robinson, M. J. F. (2013). What motivates gambling behavior? Insight into dopamine's role. *Frontiers in Behavioral Neuroscience*, 7. doi:10.3389/fnbeh.2013.00182
- Arbib, M. A. (1972). *The metaphorical brain; an introduction to cybernetics as artificial intelligence and brain theory*. Wiley-Interscience.

- Arbib, M. A., & Hesse, M. B. (1986). *The construction of reality*. Cambridge [Cambridgeshire] ; New York: Cambridge University Press.
- Artificial Brains. (2013, January 11). DARPA SyNAPSE Program. Retrieved October 24, 2014, from <http://www.artificialbrains.com/darpa-synapse-program>
- Ashby, W. R. (1960). *Design for a brain : the origin of adaptive behavior*. (2d ed.,). Chapman & Hall.
- Augstein, H. F. (1996). J C Prichard's concept of moral insanity--a medical theory of the corruption of human nature. *Medical History*, *40*, 311–343.
- Ballenger, J. C., Davidson, J. R. T., Lecrubier, Y., Nutt, D. J., Kirmayer, L. J., Lépine, J.-P., ... Ono, Y. (2001). Consensus statement on transcultural issues in depression and anxiety from the International Consensus Group on Depression and Anxiety. *Journal of Clinical Psychiatry*, *62*, 47–55.
- Bandura, A., Ross, D., & Ross, S. A. (1961). Transmission of aggression through imitation of aggressive models. *The Journal of Abnormal and Social Psychology*, *63*, 575–582.
- Barad, K. (2011). Erasers and erasures: Pinch's unfortunate "uncertainty principle." *Social Studies of Science*, *41*, 443–454.
- Barbash, F. (2015, March 27). Major publisher retracts 43 scientific papers amid wider fake peer-review scandal. *The Washington Post*. Retrieved from <http://www.washingtonpost.com/news/morning-mix/wp/2015/03/27/fabricated-peer-reviews-prompt-scientific-journal-to-retract-43-papers-systematic-scheme-may-affect-other-journals/>
- Barlett, C., Rodeheffer, C. D., Baldassarro, R., Hinkin, M. P., & Harris, R. J. (2008). The Effect of Advances in Video Game Technology and Content on Aggressive Cognitions, Hostility, and Heart Rate. *Media Psychology*, *11*, 540–565.
- Baroness Susan Greenfield: society should wake up to harmful effects of internet. (2010, September 15). *Telegraph.co.uk*. Retrieved from <http://www.telegraph.co.uk/technology/internet/8002921/Baroness-Susan-Greenfield-society-should-wake-up-to-harmful-effects-of-internet.html>
- Bavelier, D., Green, C. S., & Dye, M. W. G. (2010). Children, Wired: For Better and for Worse. *Neuron*, *67*, 692–701.
- Beatty, M. J., McCroskey, J. C., & Heisel, A. D. (1998). Communication apprehension as temperamental expression: A Communibiological paradigm. *Communication Monographs*, *65*, 197–219.
- Beaulieu, A. (2001). Voxels in the Brain. *Social Studies of Science*, *31*, 635–680.
- Beaulieu, A. (2002). Images Are Not the (Only) Truth: Brain Mapping, Visual Knowledge, and Iconoclasm. *Science, Technology, & Human Values*, *27*, 53–86.
- Bennett, C. M., Baird, A. A., Miller, M. B., & Wolford, G. L. (2011). Neural correlates of interspecies perspective taking in the post-mortem Atlantic Salmon: An argument for proper multiple comparisons correction. *Journal of Serendipitous and Unexpected Results*, *1*, 1–5.
- Berger, P. L., & Luckmann, T. (1990). *The social construction of reality: a treatise in the sociology of knowledge*. New York: Anchor Books.
- Bergson, H. (1988). *Matter and memory*. New York: Zone Books.
- Bijker, W. E. (2007). Dikes and Dams, Thick with Politics. *Isis*, *98*, 109–123.
- Black, M. (1963). *Models and Metaphors. Studies in Language and Philosophy*. Cornell University Press.

- Blaszczak-Boxe, A. (2014, May 27). DARPA program to develop brain implants for mental disorders. Retrieved October 17, 2015, from <http://www.cbsnews.com/news/darpa-program-to-develop-brain-implants-for-mental-disorders/>
- Bloor, D. (1976). The strong programme in the sociology of knowledge. *Knowledge and Social Imagery*, 2, 3–23.
- Bloor, D. (1999). Anti-Latour. *Studies In History and Philosophy of Science Part A*, 30, 81–112.
- Boden, M. A. (2006). *Mind as machine: a history of cognitive science*. Oxford : New York: Clarendon Press ; Oxford University Press.
- Bogost, I. (2013, November 8). Hyperemployment, or the Exhausting Work of the Technology User. Retrieved November 10, 2013, from <http://www.theatlantic.com/technology/archive/2013/11/hyperemployment-or-the-exhausting-work-of-the-technology-user/281149/>
- Bohil, C. J., Alicea, B., & Biocca, F. A. (2011). Virtual reality in neuroscience research and therapy. *Nature Reviews Neuroscience*, 12, 752–762.
- Boltanski, L., & Chiapello, È. (2007). *The new spirit of capitalism*. London: Verso.
- Boring, E. G. (1946). Mind and Mechanism. *The American Journal of Psychology*, 59, 173–192.
- Bowden, B. V. (1953). *Faster Than Thought*. London: Pitman. Retrieved from <http://archive.org/details/FasterThanThought>
- Bowker, G. (1993). How to be Universal: Some Cybernetic Strategies, 1943-70. *Social Studies of Science*, 23, 107–127.
- Bowker, G., & Latour, B. (1987). A Booming Discipline Short of Discipline: (Social) Studies of Science in France. *Social Studies of Science*, 17, 715–748.
- Boykoff, M. T., & Mansfield, M. (2008). “Ye Olde Hot Aire”: reporting on human contributions to climate change in the UK tabloid press. *Environmental Research Letters*, 3, 24002.
- Bradley, C., & Bowen, M. (1941). Amphetamine (benzedrine) Therapy of Children’s Behavior Disorders*. *American Journal of Orthopsychiatry*, 11, 92–103.
- Bradley, S. D. (2007). Neural Network Simulations Support Heuristic Processing Model of Cultivation Effects. *Media Psychology*, 10, 449–469.
- Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3, 77–101.
- Broadstock, M., Doughty, C., & Eggleston, M. (2007). Systematic review of the effectiveness of pharmacological treatments for adolescents and adults with autism spectrum disorder. *Autism*, 11, 335–348.
- Brockman, J. (2011). *Is the Internet changing the way you think?: the net’s impact on our minds and future*. New York: Harper Perennial.
- Bromley, E. (2006). Stimulating a normal adjustment: misbehavior, amphetamines, and the electroencephalogram at the Bradley Home for Children. *Journal of the History of the Behavioral Sciences*, 42, 379–398.
- Brooks, R. A. (1991). Intelligence without representation. *Artificial Intelligence*, 47, 139–159.
- Brown, A. (2011, January 7). What is happening in our society that we need to drug five-year-old children? *The Telegraph*. London, England.
- Brown, N. J. L., Sokal, A. D., & Friedman, H. L. (2013). The Complex Dynamics of Wishful Thinking: The Critical Positivity Ratio. *American Psychologist*. doi:10.1037/a0032850
- Brown, P. (1987). Popular Epidemiology: Community Response to Toxic Waste-Induced Disease in Woburn, Massachusetts. *Science, Technology, & Human Values*, 12, 78–85.

- Brown, P. (1995). Naming and Framing: The Social Construction of Diagnosis and Illness. *Journal of Health and Social Behavior*, 35, 34–52.
- Brown, T. (2013, December 16). An Open Letter to the Editor of the New York Times. Retrieved from <http://www.drthomasebrown.com/an-open-letter-to-the-editor-of-the-new-york-times-from-thomas-e-brown-ph-d/>
- Bruner, J., & Feldman, C. (1990). Metaphors of consciousness and cognition. In D. E. Leary (Ed.), *Metaphors in the history of psychology* (pp. 230–238). Cambridge; New York: Cambridge University Press.
- Brunshaw, J. (2011, November 15). Paying attention to the signs of ADD; Current research suggests children don't always "outgrow" disorder. *National Post*, p. AL4. Toronto.
- Bucchi, M. (1998). *Science and the media : alternative routes in scientific communication*. New York: Routledge.
- Buckels, E. E., Trapnell, P. D., & Paulhus, D. L. (2014). Trolls just want to have fun. *Personality and Individual Differences*. doi:10.1016/j.paid.2014.01.016
- Buhrmester, M., Kwang, T., & Gosling, S. D. (2011). Amazon's Mechanical Turk A New Source of Inexpensive, Yet High-Quality, Data? *Perspectives on Psychological Science*, 6, 3–5.
- Burnett, D. (2013, April 9). The only Susan Greenfield article you'll ever need. Retrieved April 9, 2013, from <http://www.guardian.co.uk/science/brain-flapping/2013/apr/09/susan-greenfield-article-how-to-guide>
- Bush, G. H. W. (1990, July 17). Decade of the Brain: Presidential Proclamation 6158 (Library of Congress). Retrieved October 22, 2014, from <http://www.loc.gov/loc/brain/proclaim.html>
- Butler, J. (1993). *Bodies That Matter: On the Discursive Limits of "Sex"*. Routledge Chapman & Hall.
- Byun, S., Ruffini, C., Mills, J. E., Douglas, A. C., Niang, M., Stepchenkova, S., ... Blanton, M. (2009). Internet Addiction: Metasynthesis of 1996–2006 Quantitative Research. *CyberPsychology & Behavior*, 12, 203–207.
- Callon, M. (1986). Some elements of a sociology of translation: domestication of the scallops and the fishermen of St Briec Bay. In J. Law (Ed.), *Power action and belief: a new sociology of knowledge?* (pp. 196–223). London: Routledge.
- Campbell, M., Hoane Jr., A. J., & Hsu, F. (2002). Deep Blue. *Artificial Intelligence*, 134, 57–83.
- Canguilhem, G. (1991). *The Normal and the Pathological*. (C. R. Fawcett, Trans.). New York: Zone.
- Carey, J. W. (1967). Harold Adams Innis and Marshall McLuhan. *The Antioch Review*, 27, 5–39.
- Carey, J. W. (2008). *Communication as Culture, Revised Edition: Essays on Media and Society*. London: Routledge.
- Carnagey, N. L., Anderson, C. A., & Bushman, B. J. (2007). The effect of video game violence on physiological desensitization to real-life violence. *Journal of Experimental Social Psychology*, 43, 489–496.
- Carr, N. (2008, August). Is Google Making Us Stupid? *The Atlantic*. Retrieved from <http://www.theatlantic.com/magazine/archive/2008/07/is-google-making-us-stupid/6868/>
- Carr, N. (2011). *The shallows: what the Internet is doing to our brains*. New York: W.W. Norton.
- Carr, N. (2013, November). All Can Be Lost: The Risk of Putting Our Knowledge in the Hands of Machines. *The Atlantic*. Retrieved from <http://www.theatlantic.com/magazine/archive/2013/11/the-great-forgetting/309516/>
- Cascio, J. (2009, August). Get Smarter. *The Atlantic*. Retrieved from <http://www.theatlantic.com/magazine/archive/2009/07/get-smarter/307548/>
- Casper, S. T. (2010). Book Review: Mind as Machine. *Medical History*, 54, 143–144.

- Changeux, J.-P. (1997). *Neuronal man: the biology of mind*. Princeton, N.J: Princeton University Press.
- Chemero, A. (2009). *Radical Embodied Cognitive Science*. The MIT Press.
- Children with ADHD; Apparent vast under-diagnosis is alarming. (2008, April 26). *Glasgow Herald*, p. 14. Glasgow, Scotland.
- Chomsky, N. (2007). Symposium on Margaret Boden, Mind as Machine: A History of Cognitive Science. *Artificial Intelligence*, 171, 1094–1103.
- Choudhury, S., & McKinney, K. A. (2013). Digital media, the developing brain and the interpretive plasticity of neuroplasticity. *Transcultural Psychiatry*, 50, 192–215.
- Choudhury, S., Nagel, S. K., & Slaby, J. (2009). Critical Neuroscience: Linking Neuroscience and Society through Critical Practice. *BioSocieties*, 4, 61–77.
- Christakis, D. A. (2014). Interactive media use at younger than the age of 2 years: Time to rethink the american academy of pediatrics guideline? *JAMA Pediatrics*, 168, 399–400.
- Christakis, D. A., Ramirez, J. S. B., & Ramirez, J. M. (2012). Overstimulation of newborn mice leads to behavioral differences and deficits in cognitive performance. *Scientific Reports*, 2. doi:10.1038/srep00546
- Christakis, D. A., Zimmerman, F. J., DiGiuseppe, D. L., & McCarty, C. A. (2004). Early Television Exposure and Subsequent Attentional Problems in Children. *Pediatrics*, 113, 708–713.
- Christopher, L. (2007, February 1). Prescription for Life? *Sydney Morning Herald*, p. 6. Sydney, Aus.
- Churchland, P. M. (1981). Eliminative Materialism and the Propositional Attitudes. *The Journal of Philosophy*, 78, 67–90.
- Clark, A. (1998). *Being There: Putting Brain, Body, and World Together Again* (1st ed.). A Bradford Book.
- Clark, A. (2003). *Natural-born cyborgs: minds, technologies, and the future of human intelligence*. Oxford; New York: Oxford University Press.
- Clarke, J. (2011). Magazine portrayal of attention deficit/hyperactivity disorder (ADD/ADHD): A post-modern epidemic in a post-trust society. *Health, Risk & Society*, 13, 621–636.
- Clarke, J. (2012). Representations of Autism in US Magazines for Women in Comparison to the General Audience. *Journal of Children and Media*, 6, 182–197.
- Clayton, R. B., Osborne, R. E., Miller, B. K., & Oberle, C. D. (2013). Loneliness, anxiousness, and substance use as predictors of Facebook use. *Computers in Human Behavior*, 29, 687–693.
- Clynes, M., & Kline, N. (1960). Cyborgs and Space. *Astronautics*, 26–27, 74–76.
- Coccaro, E. F., McCloskey, M. S., Fitzgerald, D. A., & Phan, K. L. (2007). Amygdala and Orbitofrontal Reactivity to Social Threat in Individuals with Impulsive Aggression. *Biological Psychiatry*, 62, 168–178.
- Cognitive Training Data. (2015). Retrieved February 6, 2015, from <http://www.cognitivetrainingdata.org/>
- Cohen, J. (1954). On the Project of a Universal Character. *Mind*, 63, 49–63.
- Cohen-Cole, J. (2005). The reflexivity of cognitive science: the scientist as model of human nature. *History of the Human Sciences*, 18, 107–139.
- Cohen-Cole, J. (2007). Instituting the science of mind: intellectual economies and disciplinary exchange at Harvard's Center for Cognitive Studies. *The British Journal for the History of Science*, 40. doi:10.1017/S0007087407000283
- Cohen-Cole, J. (2014). *The open mind: Cold War politics and the sciences of human nature*. Chicago ; London: The University of Chicago Press.
- Colby, K. M., Weber, S., & Hilf, F. D. (1971). Artificial Paranoia. *Artificial Intelligence*, 2, 1–25.

- Coleman, D., & Fraser, H. (Eds.). (2011). *Minds, bodies, machines, 1770 - 1930*. Basingstoke: Palgrave Macmillan.
- Collins, H. M. (1981). Son of Seven Sexes: The Social Destruction of a Physical Phenomenon. *Social Studies of Science*, 11, 33–62.
- Collins, H. M. (1983). The Sociology of Scientific Knowledge: Studies of Contemporary Science. *Annual Review of Sociology*, 9, 265–285.
- Collins, H. M. (1990). *Artificial experts: social knowledge and intelligent machines*. Cambridge, Mass.: MIT Press.
- Combrisson, E., & Jerbi, K. (2015). Exceeding chance level by chance: The caveat of theoretical chance levels in brain signal classification and statistical assessment of decoding accuracy. *Journal of Neuroscience Methods*. doi:10.1016/j.jneumeth.2015.01.010
- Comstock, E. J. (2011). The end of drugging children: Toward the genealogy of the ADHD subject. *Journal of the History of the Behavioral Sciences*, 47, 44–69.
- Conrad, P. (2005). The Shifting Engines of Medicalization. *Journal of Health and Social Behavior*, 46, 3–14.
- Conway, F., & Siegelman, J. (2009). *Dark Hero of the Information Age: In Search of Norbert Wiener, the Father of Cybernetics*. New York: Basic Books.
- Cordiano, J., & Hurst, A. Learning-Impaired Youth, § Standing Committee on Public Accounts (1994). Toronto, Can. Retrieved from http://www.ontla.on.ca/web/committee-proceedings/committee_transcripts_details.do?locale=en&Date=1994-11-03&ParlCommID=509&BillID=&Business=Learning-Impaired+Youth&DocumentID=18203
- Crary, J. (1999). *Suspensions of perception: attention, spectacle, and modern culture*. Cambridge, Mass.: MIT Press. Retrieved from <http://site.ebrary.com/id/10229581>
- Crevier, D. (1993). *AI: the tumultuous history of the search for artificial intelligence*. New York, NY: Basic Books.
- Crichton, S. A. (1798). *An Inquiry Into the Nature and Origin of Mental Derangement: Comprehending a Concise System of the Physiology and Pathology of the Human Mind. And a History of the Passions and Their Effects*. T. Cadell, Junior, and W. Davies.
- Croteau, D., & Hoynes, W. (2006). *The business of media: corporate media and the public interest* (2nd ed). Thousand Oaks, Calif: Pine Forge Press.
- Crowther-Heyck, H. (2005). *Herbert A. Simon: The Bounds of Reason in Modern America*. Baltimore: Johns Hopkins University Press.
- Crump, M. J. C., McDonnell, J. V., & Gureckis, T. M. (2013). Evaluating Amazon's Mechanical Turk as a Tool for Experimental Behavioral Research. *PLoS ONE*, 8, e57410.
- Danziger, K. (1990). Generative Metaphor and the History of Psychological Discourse. In D. E. Leary (Ed.), *Metaphors in the history of psychology* (pp. 331–355). Cambridge: Cambridge University Press.
- Danziger, K. (2008). *Marking the mind: a history of memory*. Cambridge: Cambridge University Press.
- Defense Sciences Office. (2013, October 24). Broad Agency Announcement: Systems-Based Neurotechnology for Emerging Therapies (SUBNETS). Defense Advanced Research Projects Agency. Retrieved from <https://www.fbo.gov/spg/ODA/DARPA/CMO/DARPA-BAA-14-09/listing.html>
- Deleuze, G. (1992). Postscript on the Societies of Control. *October*, 59, 3–7.

- Deleuze, G., & Guattari, F. (1987). *A thousand plateaus: capitalism and schizophrenia*. Minneapolis: University of Minnesota Press.
- Dennett, D. C. (1989). *The Intentional Stance*. Cambridge, Mass.: MIT Press.
- Derrida, J. (1981). *Dissemination*. (B. Johnson, Trans.). Chicago: University of Chicago Press.
- Descartes, R. (2003). *Treatise of man*. (T. S. Hall, Trans.). Amherst, N.Y: Prometheus Books.
- Deterding, S., Sicart, M., Nacke, L., O'Hara, K., & Dixon, D. (2011). Gamification. Using Game-design Elements in Non-gaming Contexts. In *CHI '11 Extended Abstracts on Human Factors in Computing Systems* (pp. 2425–2428). New York, NY, USA: ACM.
- DeWolf, T., & Stewart, T. (2012, December 4). We are the computational neuroscientists behind the world's largest functional brain model • /r/IAMa. Retrieved October 24, 2014, from http://www.reddit.com/r/IAMa/comments/147gqm/we_are_the_computational_neuroscientists_behind/
- Douglas, V. I. (1972). Stop, look and listen: The problem of sustained attention and impulse control in hyperactive and normal children. *Canadian Journal of Behavioural Science/Revue Canadienne Des Sciences Du Comportement*, 4, 259–282.
- Dreyfus, H. L. (1979). *What computers can't do: the limits of artificial intelligence*. New York: Harper & Row.
- Dumit, J. (2003). *Picturing Personhood: Brain Scans and Biomedical Identity* (First Paperback Edition). Princeton University Press.
- Dumit, J. (2012). *Drugs for life: how pharmaceutical companies define our health*. Durham, NC: Duke University Press.
- Dupuy, J.-P. (2000). *The Mechanization of the Mind*. (M. B. DeBevoise, Trans.). Princeton University Press.
- Dyson, E. (2012, December 21). The Rise of the Attention Economy. Retrieved November 22, 2013, from <http://www.project-syndicate.org/commentary/the-rise-of-the-attention-economy-by-esther-dyson>
- Dyson, G. (2012). *Turing's Cathedral: The Origins of the Digital Universe* (Digital). New York: Pantheon.
- Eagleman, D. (2011). *Incognito: The Secret Lives of The Brain*. Edinburgh: Canongate Books.
- Ebaugh, F. G. (1923). Neuropsychiatric sequelae of acute epidemic encephalitis in children. *American Journal of Diseases of Children*, 25, 89–97.
- Edge, D. (1974). Technological Metaphor and Social Control. *New Literary History*, 6, 135.
- Edwards, P. N. (1997). *The Closed World: Computers and the Politics of Discourse in Cold War America*. Cambridge, Mass.: The MIT Press.
- Egan, T. (2016, January 22). The Eight-Second Attention Span. *The New York Times*. Retrieved from <http://www.nytimes.com/2016/01/22/opinion/the-eight-second-attention-span.html>
- Eisenberg, L. (1986). Mindlessness and brainlessness in psychiatry. *The British Journal of Psychiatry*, 148, 497–508.
- Eisenstein, E. L. (1979). *The printing press as an agent of change: communications and cultural transformations in early modern Europe, volumes I and II*. Cambridge [England]; New York: Cambridge University Press.
- Ekbja, H. R. (2008). *Artificial Dreams: The Quest for Non-Biological Intelligence* (1st ed.). Cambridge University Press.
- Eklund, A., Nichols, T. E., & Knutsson, H. (2016). Cluster failure: Why fMRI inferences for spatial extent have inflated false-positive rates. *Proceedings of the National Academy of Sciences*, 113, 7900–7905.
- Elias, M. (2005, March 31). So much media, so little attention span. *Washington Post*, p. 6D.

- Eliasmith, C. (2013). *How to Build a Brain: A Neural Architecture for Biological Cognition* (1 edition). Oxford: Oxford University Press.
- Eliasmith, C., Stewart, T. C., Choo, X., Bekolay, T., DeWolf, T., Tang, Y., & Rasmussen, D. (2012). A Large-Scale Model of the Functioning Brain. *Science*, *338*, 1202–1205.
- Eliasmith, C., & Trujillo, O. (2014). The use and abuse of large-scale brain models. *Current Opinion in Neurobiology*, *25*, 1–6.
- Ellul, J. (1964). *The technological society*. New York: Vintage Books.
- Epstein, R. (2016, May 18). Your brain does not process information and it is not a computer. Retrieved May 22, 2016, from <https://aeon.co/essays/your-brain-does-not-process-information-and-it-is-not-a-computer>
- Epstein, S. (1998). *Impure Science: AIDS, Activism, and the Politics of Knowledge* (Reprint). University of California Press.
- Erman, L. D., Hayes-Roth, F., Lesser, V. R., & Reddy, D. R. (1980). The Hearsay-II Speech-Understanding System: Integrating Knowledge to Resolve Uncertainty. *ACM Comput. Surv.*, *12*, 213–253.
- European Commission. (2014). FET Flagships. Retrieved October 24, 2014, from ec.europa.eu//digital-agenda/en/fet-flagships
- Eyal, G. (2010). *The Autism Matrix* (1st ed.). Polity.
- Feigenbaum, E. A. (1961). Soviet Cybernetics and Computer Sciences, 1960. *Commun. ACM*, *4*, 566–579.
- Feingold, B. F. (1975). *Why your child is hyperactive* ([1st ed.]). New York: Random House.
- Ferguson, C. J., & Kilburn, J. (2010). Much ado about nothing: The misestimation and overinterpretation of violent video game effects in Eastern and Western nations: Comment on Anderson et al. (2010). *Psychological Bulletin*, *136*, 174–178.
- Fitzgerald, D., Matusall, S., Skewes, J., & Roepstorff, A. (2014). What's so critical about Critical Neuroscience? Rethinking experiment, enacting critique. *Frontiers in Human Neuroscience*, *8*. doi:10.3389/fnhum.2014.00365
- Flusser, V. (2002). *Writings*. (E. Eisel, Trans., A. Strohl, Ed.). Minneapolis, MN: University of Minnesota Press.
- Fodor, J. A. (2008). *LOT 2: The language of thought revisited*. Oxford: Oxford University Press.
- Fodor, J. A., & Pylyshyn, Z. W. (1988). Connectionism and cognitive architecture: A critical analysis. *Cognition*, *28*, 3–71.
- Foucault, M. (1990). *The History of Sexuality, Vol. 1: An Introduction*. (R. Hurley, Trans.) (Fifth Edition). New York: Vintage.
- Foucault, M. (1995). *Discipline and punish: the birth of the prison* (2nd Vintage Books ed). New York: Vintage Books.
- Foucault, M. (2003). *Abnormal: Lectures at the Collège de France, 1974-1975*. (G. Burchell, Trans.) (Reprint). London: Picador.
- Foucault, M. (2006). *Psychiatric Power: Lectures at the Collège de France, 1973--1974*. (G. Burchell, Trans., J. Lagrange & A. I. I. Davidson, Eds.) (Reprint). London: Picador.
- Freud, S. (1999). *The Standard Edition of the Complete Psychological Works of Sigmund Freud*. (J. Strachey, Ed.) (Vols. 1–24). New York: W.W. Norton.
- Fuller, S. (1989). Of Conceptual Intersections and Verbal Collisions: Towards the Routing of Slezak. *Social Studies of Science*, *19*, 625–638.
- Galison, P. (1994). The ontology of the enemy: Norbert Wiener and the cybernetic vision. *Critical Inquiry*, *21*, 228–266.

- Galison, P. (1997). *Image and logic: a material culture of microphysics*. Chicago: University of Chicago Press.
- Gardner, H. E. (2008). *The Mind's New Science: a History Of The Cognitive Revolution*. New York: Basic Books. Retrieved from <http://public.eblib.com/EBLPublic/PublicView.do?ptiID=978629>
- Gee, J. P. (2005). *Why Video Games are Good for Your Soul*. Melbourne, Vic.: Common Ground.
- Gentner, D., Holyoak, K. J., & Kokinov, B. N. (Eds.). (2001). *The Analogical Mind: Perspectives from Cognitive Science*. Cambridge, Mass.: Bradford.
- Gerbner, G. (1998). Cultivation Analysis: An Overview. *Mass Communication and Society*, 1, 175–194.
- Gerbner, G., & Gross, L. (1976). Living With Television: The Violence Profile. *Journal of Communication*, 26, 172–194.
- Gerbner, G., Gross, L., Jackson-Beeck, M., Jeffries-Fox, S., & Signorielli, N. (1978). Cultural Indicators: Violence Profile No. 9. *Journal of Communication*, 28, 176–207.
- Gerbner, G., Gross, L., Morgan, M., & Signorielli, N. (1980). The “Mainstreaming” of America: Violence Profile No. 11. *Journal of Communication*, 30, 10–29.
- Gerbner, G., Gross, L., Signorielli, N., Morgan, M., & Jackson-Beeck, M. (1979). The Demonstration of Power: Violence Profile No. 10. *Journal of Communication*, 29, 177–196.
- Gerovitch, S. (2001). “Russian Scandals”: Soviet Readings of American Cybernetics in the Early Years of the Cold War. *The Russian Review*, 60, 545–568.
- Gerovitch, S. (2002). Love-Hate for Man-Machine Metaphors in Soviet Physiology: From Pavlov to “Physiological Cybernetics.” *Science in Context*, 15, 339–374.
- Geuss, M. (2014, June 4). Why eye tracking could make VR displays like the Oculus Rift consumer-ready. Retrieved April 14, 2015, from <http://arstechnica.com/gaming/2014/06/why-eye-tracking-could-make-vr-displays-like-the-oculus-rift-consumer-ready/>
- Giedd, J. N., Blumenthal, J., Molloy, E., & Castellanos, F. X. (2001). Brain Imaging of Attention Deficit/Hyperactivity Disorder. *Annals of the New York Academy of Sciences*, 931, 33–49.
- Gieryn, T. F. (1983). Boundary-Work and the Demarcation of Science from Non-Science: Strains and Interests in Professional Ideologies of Scientists. *American Sociological Review*, 48, 781–795.
- Gigerenzer, G. (1996). From tools to theories: Discovery in cognitive psychology. *Historical Dimensions of Psychological Discourse*, 36.
- Gigerenzer, G. (2008). Why Heuristics Work. *Perspectives on Psychological Science*, 3, 20–29.
- Gleick, J. (2011). *The information: a history, a theory, a flood* (1st Vintage Books digital ed.). New York: Vintage Books.
- Glusker, A. (1997, March 30). Deficit Selling; In the strange new world of attention deficit disorder, science and medicine have mingled so thoroughly with capitalism and 12-step revivalism that it's hard to tell where the boundaries are anymore. *Washington Post*, p. W12. Washington, D.C.
- Goldacre, B. (2012). *Bad Pharma: How Drug Companies Mislead Doctors and Harm Patients*. London: Faber & Faber.
- Goldstein, D. G., & Gigerenzer, G. (2002). Models of ecological rationality: The recognition heuristic. *Psychological Review*, 109, 75–90.
- Gorman, M. (1989). Beyond Strong Programmes: How Cognitive Approaches Can Complement SSK. *Social Studies of Science*, 19, 643–653.
- Gray, J. S. (1935). An objective theory of emotion. *Psychological Review*, 42, 108–116.
- Green, C. S., & Bavelier, D. (2003). Action video game modifies visual selective attention. *Nature*, 423, 534–537.
- Greenfield, P. M. (1984). *Mind and Media*. Cambridge, MA, USA: Harvard University Press.

- Greenfield, P. M., & Bruner, J. (1966). Culture and Cognitive Growth. *International Journal of Psychology*, 1, 89–107.
- Greenfield, P., & Yan, Z. (2006). Children, Adolescents, and the Internet: A New Field of Inquiry in Developmental Psychology. *Developmental Psychology*, 42, 391–394.
- Greenfield, S. (2013, April 6). “Facebook Home could change our brains.” Retrieved from <http://www.telegraph.co.uk/technology/facebook/9975118/Facebook-Home-could-change-our-brains.html>
- Greenfield, S. (2014). *Mind change: how digital technologies are leaving their mark on our brains*. London: Rider Books.
- Hacking, I. (1986). Making up people. *Reconstructing Individualism: Autonomy, Individuality, and the Self in Western Thought*, 222–36.
- Hacking, I. (2000). *The Social Construction of What?* Cambridge, Mass.: Harvard University Press.
- Hallowell, E. (2012, January 31). Dr. Hallowell’s Response to NY Times Piece “Ritalin Gone Wrong.” Retrieved September 3, 2013, from <http://www.drhallowell.com/blog/dr-hallowells-response-to-ny-times-piece-ritalin-gone-wrong/>
- Hallowell, E. M., & Ratey, J. J. (2011). *Driven to distraction: recognizing and coping with attention deficit disorder from childhood through adulthood*. New York, NY: Anchor Books.
- Haraway, D. J. (1990). *Simians, Cyborgs, and Women: The Reinvention of Nature* (1st ed.). London: Routledge.
- Haraway, D. J. (2004). *The Haraway reader*. London: Routledge.
- Haraway, D. J. (2007). *When Species Meet*. Minneapolis: University of Minnesota Press.
- Hardaker, C. (2013, August 3). What is turning so many young men into internet trolls? Retrieved December 18, 2014, from <http://www.theguardian.com/media/2013/aug/03/how-to-stop-trolls-social-media>
- Harman, G. (2009). *Prince of Networks: Bruno Latour and Metaphysics*. Victoria, Australia: Re.Press.
- Harpy: *Speech Recognition 1976 CMU*. (1971). Retrieved from http://www.youtube.com/watch?v=N3i6NoUZsSw&feature=youtube_gdata_player
- Harrington, N. G., Lane, D., Donohew, L., & Zimmerman, R. (2006). An Extension of the Activation Model of Information Exposure: The Addition of a Cognitive Variable to a Model of Attention. *Media Psychology*, 8, 139–164.
- Harris, G., & Carey, B. (2008, June 8). Researchers Fail to Reveal Full Drug Pay. *The New York Times*. Retrieved from <http://www.nytimes.com/2008/06/08/us/08conflict.html>
- Harris, S. (2012). *Free will* (1st Free Press trade pbk. ed). New York: Free Press.
- Haugeland, J. (1989). *Artificial intelligence: the very idea*. Cambridge, Mass.: MIT Press.
- Hawkins, R. P., & Pingree, S. (1981). Uniform Messages and Habitual Viewing: Unnecessary Assumptions in Social Reality Effects. *Human Communication Research*, 7, 291–301.
- Hay, J. C., Lynch, B. E., Smith, D. R., & Murray, A. E. (1960). *Mark 1 Perceptron Operators’ Manual (Project PARA)*. (F. Rosenblatt, Ed.). Buffalo, New York: Cornell Aeronautical Laboratory.
- Hayles, K. (1994). Boundary Disputes: Homeostasis, Reflexivity, and the Foundations of Cybernetics. *Configurations*, 2, 441–467.
- Hayles, K. (1999). *How we became posthuman: virtual bodies in cybernetics, literature, and informatics*. Chicago: University of Chicago Press.
- Hayles, K. (2007). Hyper and deep attention: The generational divide in cognitive modes. *Profession*, 187–199.
- Haynes, J.-D., Sakai, K., Rees, G., Gilbert, S., Frith, C., & Passingham, R. E. (2007). Reading Hidden Intentions in the Human Brain. *Current Biology*, 17, 323–328.

- Hayward, R. (2001). The Tortoise and the Love-Machine: Grey Walter and the Politics of Electroencephalography. *Science in Context*, 14, 615–641.
- Healy, D. (2002). *The Creation of Psychopharmacology* (1st ed.). Cambridge, Mass.: Harvard University Press.
- Healy, D. (2012, November 2). Benefit risk madness: antipsychotics and suicide. Retrieved August 26, 2013, from <http://davidhealy.org/benefit-risk-madness-antipsychotics-and-suicide/>
- Heidegger, M. (1976). Only a God can save Us: Der Spiegel's Interview with Heidegger. *Philosophy Today*, 20. Retrieved from <http://search.proquest.com.ezproxy.library.yorku.ca/pao/docview/1301479331/citation?accountid=15182>
- Heidegger, M. (1977). *The question concerning technology, and other essays*. New York: Garland Pub.
- Heims, S. J. (1991). *The cybernetics group*. Cambridge, Mass.: MIT Press.
- Heims, S. J. (1993). *Constructing a social science for postwar America: the cybernetics group, 1946-1953*. Cambridge, Mass: MIT Press.
- Hendler, J. (2008). A Letter from the Editor: Avoiding Another AI Winter. *IEEE Intelligent Systems*, 2–4.
- Heyck, H. (2008a). Defining the Computer: Herbert Simon and the Bureaucratic Mind—Part 1. *IEEE Annals of the History of Computing*, 1, 42–51.
- Heyck, H. (2008b). Defining the Computer: Herbert Simon and the Bureaucratic Mind—Part 2. *IEEE Annals of the History of Computing*, 1, 52–63.
- Heyck, H. (2014). The Organizational Revolution and the Human Sciences. *Isis*, 105, 1–31.
- Hicks, S. (2014, February 6). Does the Video Game Industry Hold the Keys to the Future of Advertising? Retrieved December 16, 2014, from <http://www.adweek.com/news/advertising-branding/engage-consumers-brain-chemistry-155531>
- Hinton, G. (2007). Boltzmann machine. *Scholarpedia*, 2, 1668.
- Hobbes, T. (1651). *Leviathan*. London: Andrew Crooke.
- Hofstadter, D. R. (1979). *Gödel, Escher, Bach: An Eternal Golden Braid*. New York: Basic Books.
- Hofstadter, D. R. (1996). *Fluid Concepts And Creative Analogies: Computer Models Of The Fundamental Mechanisms Of Thought* (1st ed.). Basic Books.
- Holland, O. (2003). Exploration and high adventure: the legacy of Grey Walter. *Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, 361, 2085–2121.
- Horvath, J. C. (2015, February 3). Does brain training work? That depends on your purpose. Retrieved February 6, 2015, from <http://theconversation.com/does-brain-training-work-that-depends-on-your-purpose-36947>
- Hsu, J. (2015, January 27). Inside the Largest Virtual Psychology Lab in the World. Retrieved February 20, 2015, from <https://medium.com/backchannel/inside-the-largest-virtual-psychology-lab-in-the-world-7c0d2c43cda5>
- Huang, C. (2010). Internet use and psychological well-being: a meta-analysis. *Cyberpsychology, Behavior and Social Networking*, 13, 241–249.
- Hughes, A. C., & Hughes, T. P. (Eds.). (2000). *Systems, experts, and computers: the systems approach in management and engineering, World War II and after*. Cambridge, Mass: MIT Press.
- Hughes, W. (2014, June 27). Facebook tinkered with users' feeds for a massive psychology experiment. Retrieved March 27, 2015, from <http://www.avclub.com/article/facebook-tinkered-users-feeds-massive-psychology-e-206324>

- Human Brain Project. (2012). Project Objectives - Human Brain Project. Retrieved October 22, 2014, from <https://www.humanbrainproject.eu/project-objectives>
- Husbands, P., Holland, O., & Wheeler, M. (2008). *The mechanical mind in history*. Cambridge, Mass.; London: MIT. Retrieved from <http://site.ebrary.com/id/10214161>
- Hutchins, E. (1996). *Cognition in the Wild* (1st ed.). Cambridge, Mass.: A Bradford Book.
- Hutchins, E. (2014). The cultural ecosystem of human cognition. *Philosophical Psychology*, 27, 34–49.
- Hutton, N. (2012). *Frontiers | Bluebrain | Year Three*. Retrieved from http://www.frontiersin.org/video/Bluebrain_%7C_Year_Three/1775
- IBM Research. (2013, December 11). Why cognitive systems? Retrieved October 24, 2014, from <http://www.research.ibm.com/cognitive-computing/why-cognitive-systems.shtml>
- Innis, H. A. (2007). *Empire and communications*. (A. J. Watson, Ed.). Toronto: Dundurn Press.
- Ivory, J. D., & Kalyanaraman, S. (2007). The Effects of Technological Advancement and Violent Content in Video Games on Players' Feelings of Presence, Involvement, Physiological Arousal, and Aggression. *Journal of Communication*, 57, 532–555.
- Jennings, K. (2011, February 16). My Puny Human Brain. *Slate*. Retrieved from http://www.slate.com/articles/arts/culturebox/2011/02/my_puny_human_brain.html
- Johnson, J. (1988). Mixing Humans and Nonhumans Together: The Sociology of a Door-Closer. *Social Problems*, 35, 298–310.
- Johnson, S. (2005). *Everything bad is good for you: how today's popular culture is actually making us smarter*. New York: Riverhead Books.
- Jones, E. G. (1999). Assessing the Decade of the Brain. *Science*, 284, 739–739.
- Jones, S. C., & Harwood, V. (2009). Representations of autism in Australian print media. *Disability & Society*, 24, 5–18.
- Kahlig, K. M., Binda, F., Khoshbouei, H., Blakely, R. D., McMahon, D. G., Javitch, J. A., & Galli, A. (2005). Amphetamine induces dopamine efflux through a dopamine transporter channel. *Proceedings of the National Academy of Sciences of the United States of America*, 102, 3495–3500.
- Kahneman, D. (2011). *Thinking, Fast and Slow* (1st ed.). Farrar, Straus and Giroux.
- Kang, S. (2013). Coverage of autism spectrum disorder in the US television news: an analysis of framing. *Disability & Society*, 28, 245–259.
- Kanner, L. (1972). *Child psychiatry* (4th ed.). Springfield, Ill.: C. C. Thomas.
- Katz, E. (1999). Foreword: Two Virtual Debates Between Lazarsfeld and McLuhan On Radio. *Journal of Radio Studies*, 6, 5–9.
- Kay, L. E. (2000). *Who wrote the book of life?: a history of the genetic code*. Stanford, Calif.: Stanford University Press.
- Keith, A. (1920). *The engines of the human body*. Philadelphia: Lippincott.
- Keller, J. D., Bazerman, C., & Latour, B. (1996). Cognition in the Wild (Book Review). *Mind, Culture, and Activity*, 3, 46–63.
- Kelly, K. (2014, October 27). The Three Breakthroughs That Have Finally Unleashed AI on the World. Retrieved December 13, 2014, from <http://www.wired.com/2014/10/future-of-artificial-intelligence/>
- Kendler, K. S. (2009). An historical framework for psychiatric nosology. *Psychological Medicine*, 39, 1935–1941.
- Kendler, K. S., Muñoz, R. A., & Murphy, G. (2009). The Development of the Feighner Criteria: A Historical Perspective. *American Journal of Psychiatry*, 167, 134–142.

- Kennedy, J. L., & Chapman, R. (1955). The Background and Implications of the Systems Research Laboratory Studies. RAND Corporation. Retrieved from <http://www.rand.org/pubs/papers/P740.html>
- Keynes, J. M. (1963). *Essays in Persuasion*. New York: W. W. Norton.
- Kirmayer, L. J. (2006). Culture and Psychotherapy in a Creolizing World. *Transcultural Psychiatry*, 43, 163–168.
- Kittler, F. (1990). *Discourse networks 1800/1900*. Stanford, Calif: Stanford University Press.
- Kittler, F. (1999). *Gramophone, Film, Typewriter*. (G. Winthrop-Young & M. Wutz, Trans.) (1st ed.). Stanford University Press.
- Klass, P. (2011, May 10). Fixated by Screens, but Seemingly Nothing Else. *New York Times*, p. 5.
- Klatt, D. H. (1977). Review of the ARPA Speech Understanding Project. *The Journal of the Acoustical Society of America*, 62, 1345–1366.
- Kline, R. (2009). Where are the Cyborgs in Cybernetics? *Social Studies of Science*, 39, 331–362.
- Koepp, M. J., Gunn, R. N., Lawrence, A. D., Cunningham, V. J., Dagher, A., Jones, T., ... Grasby, P. M. (1998). Evidence for striatal dopamine release during a video game. *Nature*, 393, 266–268.
- Kosoff, M. (2014, July 10). The Incredibly Sad Story Of How One Couple Let Their Child Die Because They Were Addicted To The Internet. Retrieved December 16, 2014, from <http://www.businessinsider.com/south-koreas-internet-addiction-epidemic-2014-7>
- Kraepelin, E., & Diefendorf, A. R. (1915). *Clinical psychiatry: a text-book for students and physicians*. The Macmillan Company.
- Kramer, A. D. I., Guillory, J. E., & Hancock, J. T. (2014). Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences*, 111, 8788–8790.
- Kramer, P. D. (1997). *Listening to Prozac*. New York, N.Y.: Penguin Books.
- Kraut, R., Kiesler, S., Boneva, B., Cummings, J., Helgeson, V., & Crawford, A. (2002). Internet Paradox Revisited. *Journal of Social Issues*, 58, 49–74.
- Kraut, R., Patterson, M., Lundmark, V., Kiesler, S., Mukophadhyay, T., & Scherlis, W. (1998). Internet paradox: A social technology that reduces social involvement and psychological well-being? *American Psychologist*, 53, 1017–1031.
- Krisch, J. A. (2014, April 3). Why is the 2048 Game So Addictive? Retrieved December 16, 2014, from <http://www.popularmechanics.com/technology/gadgets/video-games/why-the-2048-game-is-so-addictive-16659899>
- Kühn, S., Romanowski, A., Schilling, C., Lorenz, R., Mörsen, C., Seiferth, N., ... Gallinat, J. (2011). The neural basis of video gaming. *Translational Psychiatry*, 1, e53.
- Kuhn, T. S. (1957). *The Copernican revolution: planetary astronomy in the development of Western thought*. Cambridge: Harvard University Press.
- Kuhn, T. S. (2012). *The structure of scientific revolutions* (Fourth edition). Chicago ; London: The University of Chicago Press.
- Kurzweil, R. (1990). *The age of intelligent machines*. Cambridge, Mass: MIT Press.
- Lakatos, I. (1968). Criticism and the Methodology of Scientific Research Programmes. *Proceedings of the Aristotelian Society*, 69, 149–186.
- Lakoff, A. (2000). Adaptive will: the evolution of attention deficit disorder. *Journal of the History of the Behavioral Sciences*, 36, 149–169.

- Lakoff, A. (2005). *Pharmaceutical Reason: Knowledge and Value in Global Psychiatry*. Cambridge University Press.
- Lakoff, G., & Johnson, M. (1999). *Philosophy in the Flesh: The Embodied Mind and Its Challenge to Western Thought*. New York, N.Y.: Basic Books.
- Lakoff, G., & Johnson, M. (2003). *Metaphors we live by*. Chicago: University of Chicago Press.
- Lang, A. (2000). The limited capacity model of mediated message processing. *Journal of Communication*, 50, 46–70.
- Lange, K. W., Reichl, S., Lange, K. M., Tucha, L., & Tucha, O. (2010). The history of attention deficit hyperactivity disorder. *Attention Deficit and Hyperactivity Disorders*, 2, 241–255.
- Langley, P. (1979). Rediscovering Physics with BACON 3. In *Proceedings of the Sixth Annual Joint Conference on Artificial Intelligence*. Tokyo, Jpn. Retrieved from <http://doi.library.cmu.edu/10.1184/pmc/simon/box00013/fld00898/bdl0001/doc0001>
- Langley, P. (1987). *Scientific Discovery: Computational Explorations of the Creative Processes*. MIT Press.
- Lashley, K. S. (1930). Basic neural mechanisms in behavior. *Psychological Review*, 37, 1–24.
- Latil, P. de. (1956). *Thinking by Machine: A Study of Cybernetics*. (M. Golla, Trans.). Boston: Houghton Mifflin.
- Latour, B. (1979). *Laboratory life: the social construction of scientific facts*. Beverly Hills: Sage Publications.
- Latour, B. (1986). Visualization and cognition. *Knowledge and Society*, 6, 1–40.
- Latour, B. (1988). *Science in Action: How to Follow Scientists and Engineers through Society* (REP). Harvard University Press.
- Latour, B. (1993a). *The Pasteurization of France*. (A. Sheridan & J. Law, Trans.) (Reprint edition). Cambridge, Mass.: Harvard University Press.
- Latour, B. (1993b). *We have never been modern*. Cambridge, Mass: Harvard University Press.
- Latour, B. (1994). On Technical Mediation - Philosophy, Sociology, Genealogy. *Common Knowledge*, 3, 29–64.
- Latour, B. (1996). Social theory and the study of computerized work sites. In W. J. Orlikowski (Ed.), *Information technology and changes in organizational work* (pp. 295–307). London: Chapman & Hall.
- Latour, B. (1999). For David Bloor... and Beyond: A Reply to David Bloor's 'Anti-Latour'. *Studies in History and Philosophy of Science Part A*, 30, 113–129.
- Latour, B. (2002). Morality and Technology: The End of the Means. *Theory, Culture & Society*, 19, 247–260.
- Latour, B. (2004a). How to Talk About the Body? the Normative Dimension of Science Studies. *Body & Society*, 10, 205–229.
- Latour, B. (2004b). Why has critique run out of steam? From matters of fact to matters of concern. *Critical Inquiry*, 30, 225–248.
- Latour, B. (2007). *Reassembling the Social: An Introduction to Actor-Network-Theory*. Oxford University Press, USA.
- Latour, B., & Weibel, P. (Eds.). (2005). *Making Things Public: Atmospheres of Democracy* (1st ed.). Cambridge, Mass.: The MIT Press.
- Law, J. (1987). Technology and heterogeneous engineering: the case of Portuguese expansion. In W. Bijker, T. P. Hughes, & T. Pinch (Eds.), *The social construction of technological systems: New directions in the sociology and history of technology* (Vol. 1, pp. 1–134).
- Leary, D. E. (Ed.). (1990). *Metaphors in the history of psychology*. Cambridge: Cambridge University Press.

- Lederberg, J. (1987, November 5). How DENDRAL Was Conceived and Born [Drafts (documents)]. Retrieved October 1, 2014, from <http://profiles.nlm.nih.gov/BB/A/L/Y/P/>
- Lehrer, J. (2011). Out of the Blue. Retrieved October 23, 2014, from http://seedmagazine.com/content/article/out_of_the_blue/
- Lenoir, T. (1986). Models and Instruments in the Development of Electrophysiology, 1845-1912. *Historical Studies in the Physical and Biological Sciences*, 17, 1–54.
- Lenoir, T. (1994). Helmholtz and the Materialities of Communication. *Osiris*, 9, 184–207.
- Levine, M., Gow, R., & Shea, S. (2005). Report of New Drug Committee on Adderall XR - Canada. Government of Canada.
- Levinson, P. (2001). *Digital McLuhan: a guide to the information millenium*. London: Routledge.
- Levy, D. T., & Friend, K. (2001). A computer simulation model of mass media interventions directed at tobacco use. *Preventive Medicine*, 32, 284–294.
- Levy, S. (2011). *In the plex: how Google thinks, works, and shapes our lives* (1st Simon & Schuster hardcover ed). New York: Simon & Schuster.
- Lewis, D. (1999). *Papers in Metaphysics and Epistemology*. Cambridge: Cambridge University Press.
- Licklider, J. C. R. (1960). Man-Computer Symbiosis. *IRE Transactions on Human Factors in Electronics*, HFE-1, 4–11.
- Licklider, J. C. R. (1988, October 28). J.C.R. Licklider Interview. Retrieved from <http://conservancy.umn.edu/bitstream/107436/1/oh150jcl.pdf>
- Licklider, J. C. R. (1990). In Memoriam: J.C. Licklider (reprinting of “Man-Computer Symbiosis” and “The Computer as a Communication Device”). (R. Taylor, Ed.). Digital Equipment Corporation.
- Light, J. S. (1999). When Computers Were Women. *Technology and Culture*, 40, 455–483.
- Lighthill, J. (1973). Artificial Intelligence: A General Survey. Science Research Council. Retrieved from http://www.chilton-computing.org.uk/inf/literature/reports/lighthill_report/contents.htm
- Linnet, J., Peterson, E., Doudet, D. J., Gjedde, A., & Møller, A. (2010). Dopamine release in ventral striatum of pathological gamblers losing money. *Acta Psychiatrica Scandinavica*, 122, 326–333.
- López-Muñoz, F., Alamo, C., Cuenca, E., Shen, W., Clervoy, P., & Rubio, G. (2005). History of the Discovery and Clinical Introduction of Chlorpromazine. *Annals of Clinical Psychiatry*, 17, 113–135.
- Lucas, C. P. (1992). The order effect: reflections on the validity of multiple test presentations. *Psychological Medicine*, 22, 197–202.
- MacKenzie, D. A. (1998). *Knowing machines: essays on technical change*. Cambridge, Mass. [u.a.: The MIT Press.
- Malabou, C. (2008). *What should we do with our brain?* New York: Fordham University Press.
- Marche, S. (2012, May). Is Facebook Making Us Lonely? *The Atlantic*. Retrieved from <http://www.theatlantic.com/magazine/archive/2012/05/is-facebook-making-us-lonely/8930/>
- Markram, H. (2009). *A brain in a supercomputer*. Retrieved from http://www.ted.com/talks/henry_markram_supercomputing_the_brain_s_secrets?language=en
- Markram, H. (2012, June). The Human Brain Project. *Scientific American*, 34–39.
- Markram, K., & Markram, H. (2010). The Intense World Theory - A Unifying Theory of the Neurobiology of Autism. *Frontiers in Human Neuroscience*, 4. doi:10.3389/fnhum.2010.00224

- Marr, D., & Poggio, T. (1976). *From Understanding Computation to Understanding Neural Circuitry*. Retrieved from <http://dspace.mit.edu/handle/1721.1/5782>
- Marshall McLuhan Interview 1967*. (1967). Retrieved from http://www.youtube.com/watch?v=OMEC_HqWIBY&feature=youtube_gdata_player
- Marx, L. (1997). "Technology": The Emergence of a Hazardous Concept. *Social Research*, 64, 965–988.
- Mathews, V. P., Kronenberger, W. G., Wang, Y., Lurito, J. T., Lowe, M. J., & Dunn, D. W. (2005). Media violence exposure and frontal lobe activation measured by functional magnetic resonance imaging in aggressive and nonaggressive adolescents. *Journal of Computer Assisted Tomography*, 29, 287–292.
- Mathiak, K., & Weber, R. (2006). Toward brain correlates of natural behavior: fMRI during violent video games. *Human Brain Mapping*, 27, 948–956.
- Maturana, H. R., & Varela, F. J. (1980). *Autopoiesis and cognition: the realization of the living*. Dordrecht, Holland ; Boston: D. Reidel Pub. Co.
- Max Planck Institute for Human Development and Stanford Center on Longevity. (2014, October 20). A Consensus on the Brain Training Industry from the Scientific Community. Retrieved February 6, 2015, from <http://longevity3.stanford.edu/blog/2014/10/15/the-consensus-on-the-brain-training-industry-from-the-scientific-community-2/>
- Mayes, R., & Rafalovich, A. (2007). Suffer the restless children: the evolution of ADHD and paediatric stimulant use, 1900-80. *History of Psychiatry*, 18, 435–457.
- McAlpine, J. (2006, September 7). IS RITALIN A COVER-UP FOR BAD PARENTS? *Daily Mail*, p. 14. London, England.
- McCann, D., Barrett, A., Cooper, A., Crumpler, D., Dalen, L., Grimshaw, K., ... Stevenson, J. (2007). Food additives and hyperactive behaviour in 3-year-old and 8/9-year-old children in the community: a randomised, double-blinded, placebo-controlled trial. *Lancet*, 370, 1560–1567.
- McCarthy, J. (1989). Oral history interview with John McCarthy. Retrieved from <http://conservancy.umn.edu/handle/11299/107476>
- McCarthy, J. (1996). LISP prehistory - Summer 1956 through Summer 1958. Retrieved September 26, 2014, from <http://www-formal.stanford.edu/jmc/history/lisp/node2.html>
- McCorduck, P. (2004). *Machines who think: a personal inquiry into the history and prospects of artificial intelligence*. Natick, Mass.: A.K. Peters.
- McCulloch, W. S. (1954). Through the Den of the Metaphysician. *The British Journal for the Philosophy of Science*, V, 18–31.
- McCulloch, W. S., & Pitts, W. (1943). A logical calculus of the ideas immanent in nervous activity. *The Bulletin of Mathematical Biophysics*, 5, 115–133.
- McLuhan, M. (1959). Myth and Mass Media. *Daedalus*, 88, 339–348.
- McLuhan, M. (1978). The Brain and the Media: The "Western" Hemisphere. *Journal of Communication*, 28, 54–60.
- McLuhan, M. (1994). *Understanding media: the extensions of man* (1st MIT Press ed.). Cambridge, Mass: MIT Press.
- McLuhan, M., & Fiore, Q. (2001). *The medium is the massage*. Corte Madera: Gingko Pr.
- McLuhan, M., & McLuhan, E. (1999). *Laws of media: the new science* (Repr). Toronto: Univ. of Toronto Press.

- McReynolds, P. (1980). The clock metaphor in the history of psychology. In T. Nickles (Ed.), *Scientific discovery: Case studies* (pp. 97–112). Berlin: Springer.
- Messaris, P., & Humphreys, L. (2006). *Digital Media: Transformations in Human Communication*. Peter Lang.
- Meyrowitz, J. (1998). Multiple media literacies. *Journal of Communication*, 48, 96–108.
- Mick, H. (2008, March 11). Tot terror; Anxiety disorders are surprisingly common among children - more prevalent than attention-deficit disorder - yet help is in short supply. *The Globe and Mail*, p. L1. Toronto, Can.
- Miller, G. A. (1953). What is information measurement? *American Psychologist*, 8, 3–11.
- Miller, G. A. (1956). The magical number seven, plus or minus two: some limits on our capacity for processing information. *Psychological Review*, 63, 81–97.
- Miller, G. A. (2003). The cognitive revolution: a historical perspective. *Trends in Cognitive Sciences*, 7, 141–144.
- Miller, G. A., Galanter, E., & Pribram, K. (1960). *Plans and the structure of behavior*. New York: Henry Holt and Co, Inc.
- Mills, K. L. (2014). Effects of Internet use on the adolescent brain: despite popular claims, experimental evidence remains scarce. *Trends in Cognitive Sciences*, 18, 385–387.
- Mindell, D. A. (1995). Automation's Finest Hour: Bell Labs and Automatic Control in World War II. *IEEE Control Systems*, 15, 72-.
- Minsky, M. (1986). *The society of mind*. New York: Simon and Schuster.
- Minsky, M., & Papert, S. (1969). *Perceptrons: an introduction to computational geometry*. Cambridge, Mass.: MIT Press.
- Mirowski, P. (2001). *Machine Dreams: Economics Becomes a Cyborg Science*. Cambridge University Press.
- Mises, R., Quemada, N., Botbol, M., Burzsteijn, C., Garrabe, J., Golse, B., ... Thevenot, J. P. (2002). French classification for child and adolescent mental disorders. *Psychopathology*, 35, 176–180.
- Mol, A. (2003). *The Body Multiple: Ontology in Medical Practice*. Duke University Press Books.
- Molinaro, M., McLuhan, C., & Toyne, W. (Eds.). (1987). *Letters of Marshall McLuhan* (1st Edition edition). Oxford: Oxford University Press.
- Moon, N. W. (2009). The amphetamine years: a study of the medical applications and extramedical consumption of psychostimulant drugs in the postwar united states, 1945-1980. Retrieved from <http://smartech.gatech.edu/handle/1853/31743>
- Moore, A., & Stilgoe, J. (2009). Experts and Anecdotes The Role of “Anecdotal Evidence” in Public Scientific Controversies. *Science, Technology & Human Values*, 34, 654–677.
- Morawski, J. G. (2005). Reflexivity and the psychologist. *History of the Human Sciences*, 18, 77–105.
- Murphy, R. (2013, June 13). Tamagotchi, FarmVille, and “Fun Pain.” Retrieved October 24, 2015, from [http://www.gamasutra.com/blogs/RosstinMurphy/20130613/194315/Tamagotchi_FarmV ille_and_quotFun_Painquot.php](http://www.gamasutra.com/blogs/RosstinMurphy/20130613/194315/Tamagotchi_FarmVille_and_quotFun_Painquot.php)
- Muskal, M. (2012, April 11). Companies belittled risks of Risperdal, slapped with huge fine. *Los Angeles Times*. Retrieved from <http://articles.latimes.com/2012/apr/11/nation/la-na-nn-risperdal-arkansas-20120411>
- Nelkin, D. (1990). Selling science. *Physics Today*, 43, 41.
- Nelkin, D. (1995). *Selling science : how the press covers science and technology* (Rev. ed.). W.H. Freeman,

- Neuroskeptic. (2014, July 11). Interview: What's wrong with the Human Brain Project? Retrieved from <http://blogs.discovermagazine.com/neuroskeptic/2014/07/11/whats-wrong-human-brain-project/>
- NEW NAVY DEVICE LEARNS BY DOING; Psychologist Shows Embryo of Computer Designed to Read and Grow Wiser. (1958, July 8). *The New York Times*. Retrieved from <http://query.nytimes.com/gst/abstract.html?res=9D01E4D8173DE53BBC4053DFB1668383649EDE>
- Newell, A. (1963, February). A Guide to the General Problem-Solver Program GPS-2-2. The RAND Corporation. Retrieved from http://bitsavers.informatik.uni-stuttgart.de/pdf/rand/ipl/RM-3337-PR_A_Guide_To_The_General_Problem-Solver_Program_GPS-2-2_Feb63.pdf
- Newell, A. (1982, January 1). *Intellectual issues in the history of artificial intelligence*. Paper 2448, Computer Science Department, Carnegie Mellon University. Retrieved from <http://repository.cmu.edu/compsci/2448>
- Newell, A. (1991, June). Oral history interview with Allen Newell. Retrieved from <http://conservancy.umn.edu/handle/11299/107544>
- Newell, A., & Shaw, C. (1957). Programming the Logic Theory Machine. Presented at the Western Joint Computer Conference, Los Angeles, California. Retrieved from http://bitsavers.trailing-edge.com/pdf/rand/ipl/P-954_Programming_The_Logic_Theory_Machine_Feb57.pdf
- Newell, A., Shaw, J. C., & Simon, H. A. (1958a). Elements of a theory of human problem solving. *Psychological Review*, 65, 151–166.
- Newell, A., Shaw, J. C., & Simon, H. A. (1958b, December 30). Report on a General Problem-Solving Program. The Rand Corporation. Retrieved from http://bitsavers.informatik.uni-stuttgart.de/pdf/rand/ipl/P-1584_Report_On_A_General_Problem-Solving_Program_Feb59.pdf
- Newell, A., & Simon, H. . (1956). The logic theory machine—A complex information processing system. *IRE Transactions on Information Theory*, 2, 61–79.
- Newell, A., & Simon, H. A. (1972). *Human problem solving*. Englewood Cliffs, N.J: Prentice-Hall.
- Newell, A., & Simon, H. A. (1976). Computer Science As Empirical Inquiry: Symbols and Search. *Commun. ACM*, 19, 113–126.
- Nielsen, J. A., Zielinski, B. A., Ferguson, M. A., Lainhart, J. E., & Anderson, J. S. (2013). An Evaluation of the Left-Brain vs. Right-Brain Hypothesis with Resting State Functional Connectivity Magnetic Resonance Imaging. *PLoS ONE*, 8, e71275.
- Nietzsche, F. W. (1873). On Truth and Lie in an Extra-Moral Sense. Retrieved July 25, 2015, from http://oregonstate.edu/instruct/phl201/modules/Philosophers/Nietzsche/Truth_and_Lie_in_an_Extra-Moral_Sense.htm
- Noë, A. (2004). *Action in perception*. Cambridge, Mass.: MIT Press.
- Nyce, J. M., & Kahn, P. (Eds.). (1991). *From Memex to Hypertext: Vannevar Bush and the Mind's Machine*. San Diego: Academic Press Professional, Inc.
- O'Connor, C., Rees, G., & Joffe, H. (2012). Neuroscience in the Public Sphere. *Neuron*, 74, 220–226.
- Offray, J. de la M. (1748). *Man a Machine*. (Marquis D'Argens, Trans.). Dublin: W. Brien.
- Ogilvie, M. (2009, May 30). Were there always hyperactive kids? Children's attention deficit disorder didn't exist until 1950s, researcher maintains. *Toronto Star*, p. A08. Toronto, Can.
- Olazaran, M. (1996). A Sociological Study of the Official History of the Perceptrons Controversy. *Social Studies of Science*, 26, 611–659.

- Ollman, B. (1977). Marx's Vision of communism: a reconstruction. *Critique*, 8, 4–41.
- Oransky, I. (2014, October 26). Scientist sues PubPeer commenters, subpoenas site for names. Retrieved from <http://retractionwatch.com/2014/10/26/scientist-sues-pubpeer-commenters-subpoenas-site-for-names/>
- Orca, S. (2009, November 30). The Race to Reverse Engineer the Human Brain - h+ Magazine. Retrieved from <http://hplusmagazine.com/2009/11/30/race-reverse-engineer-human-brain/>
- Otis, L. (2001). *Networking: Communicating With Bodies and Machines in the Nineteenth Century*. Ann Arbor: University of Michigan Press.
- Otis, L. (2002). The Metaphoric Circuit: Organic and Technological Communication in the Nineteenth Century. *Journal of the History of Ideas*, 63, 105–128.
- Owens, A. M. (2006, April 26). A world gone ADD: A psychiatrist finds symptoms of disorder just come with our busy existence. *National Post*, p. A3. Toronto, Can.
- Packer, A. M., Roska, B., & Häusser, M. (2013). Targeting neurons and photons for optogenetics. *Nature Neuroscience*, 16, 805–815.
- Perkel, J. M. (2015). Rate that journal. *Nature*, 520, 119–120.
- Petty, R. E., & Cacioppo, J. T. (1986). The Elaboration Likelihood Model of Persuasion. In L. Berkowitz (Ed.), *Advances in Experimental Social Psychology* (Vol. 19, pp. 123–205). Academic Press.
- Pfeiffer, U. J., Vogeley, K., & Schilbach, L. (2013). From gaze cueing to dual eye-tracking: Novel approaches to investigate the neural correlates of gaze in social interaction. *Neuroscience & Biobehavioral Reviews*, 37, 2516–2528.
- Pickering, A. (2002). Cybernetics and the Mangle Ashby, Beer and Pask. *Social Studies of Science*, 32, 413–437.
- Pickering, A. (2011). *The Cybernetic Brain: Sketches of Another Future*. University Of Chicago Press.
- Pickersgill, M. (2009). Between Soma and Society: Neuroscience and the Ontology of Psychopathy. *BioSocieties*, 4, 45–60.
- Pies, R. (2009). Should DSM-V Designate “Internet Addiction” a Mental Disorder? *Psychiatry (Edgmont)*, 6, 31–37.
- Pinch, T. (2011). Review Essay: Karen Barad, quantum mechanics, and the paradox of mutual exclusivity. *Social Studies of Science*, 41, 431–441.
- Pinker, S. (2002). *The blank slate: the modern denial of human nature*. New York: Penguin.
- Pinker, S. (2010, June 10). Mind Over Mass Media. *The New York Times*. Retrieved from <http://www.nytimes.com/2010/06/11/opinion/11Pinker.html>
- Plato. (n.d.). *Phaedrus*. (B. Jowett, Trans.) (Internet Classics Archive). Retrieved from <http://classics.mit.edu/Plato/phaedrus.html>
- Pooley, J. (2006a). *An Accident Of Memory: Edward Shils, Paul Lazarsfeld and the History of American Mass Communication Research*. COLUMBIA UNIVERSITY. Retrieved from <http://jeffpooley.com/pubs/DissertationPooley2006.pdf>
- Pooley, J. (2006b). Fifteen Pages that Shook the Field: Personal Influence, Edward Shils, and the Remembered History of Mass Communication Research. *The ANNALS of the American Academy of Political and Social Science*, 608, 130–156.
- Popper, K. (1978). Natural Selection and the Emergence of Mind. *Dialectica*, 32, 339–355.
- Porter, T. M. (2012). Thin Description: Surface and Depth in Science and Science Studies. *Osiris*, 27, 209–226.

- Prater, M. A., & Pancheri, C. (1999). What Teachers and Parents Should Know about Ritalin. *TEACHING Exceptional Children*, 31, 20–26.
- Pribram, K. (1990). From metaphors to models: the use of analogy in neuropsychology. In D. E. Leary (Ed.), *Metaphors in the history of psychology* (pp. 79–103). Cambridge: Cambridge University Press.
- Prigogine, I., & Stengers, I. (1984). *Order out of chaos: man's new dialogue with nature*. New York, N.Y.: Bantam.
- Pryer, W. (2004, August 26). ADHD medication linked to adult drug dependence. *The West Australian*, p. 45. Perth, Australia.
- Pryor, C. N. (1961). The Perceptron as an Adaptive Classification Device. United States Government, Naval Ordnance Laboratory.
- Pylyshyn, Z. W. (1980). Computation and cognition: Issues in the foundations of cognitive science. *Behavioral and Brain Sciences*, 3, 111–132.
- Rabinbach, A. (1992). *The human motor: energy, fatigue, and the origins of modernity*. Berkeley: University of California Press.
- Racine, E., Bar-Ilan, O., & Illes, J. (2005). fMRI in the public eye. *Nature Reviews Neuroscience*, 6, 159–164.
- Racine, E., Waldman, S., Rosenberg, J., & Illes, J. (2010). Contemporary neuroscience in the media. *Social Science & Medicine*, 71, 725–733.
- Raeburn, P. (2010, September 2). NY Times: Congratulations! Your piece is the 1 millionth story on overmedicated kids!!! Retrieved May 22, 2013, from <http://ksj.mit.edu/tracker/2010/09/ny-times-congratulations-your-piece-1-mi>
- Rafalovich, A. (2001). The conceptual history of attention deficit hyperactivity disorder: Idiocy, imbecility, encephalitis and the child deviant, 1877-1929. *Deviant Behavior*, 22, 93–115.
- Rasmussen, N. (2008). America's First Amphetamine Epidemic 1929–1971. *Journal Information*, 98. Retrieved from <http://ajph.aphapublications.org/doi/abs/10.2105/AJPH.2007.110593>
- Reas, E. (2015, January 28). FMRI under the Microscope: An interview with MRI Physicist Practical fMRI. Retrieved from <http://neuro.plos.org/2015/01/28/fmri-under-the-microscope-an-interview-with-mri-physicist-practical-fmri/>
- Requarth, T. (2013, March 18). Bringing a Virtual Brain to Life. *The New York Times*. Retrieved from <http://www.nytimes.com/2013/03/19/science/bringing-a-virtual-brain-to-life.html>
- Rheinberger, H.-J. (1997). *Toward a history of epistemic things: synthesizing proteins in the test tube*. Stanford, Calif: Stanford University Press.
- Richtel, M. (2003, July 6). The Lure of Data: Is It Addictive? *New York Times*, p. 1.
- Ricœur, P. (1993). *The rule of metaphor multi-disciplinary studies of the creation of meaning in language*. Toronto: University of Toronto Press. Retrieved from <http://search.ebscohost.com/login.aspx?direct=true&scope=site&db=nlebk&db=nlabk&AN=468632>
- Roberts, E. (2006, October 8). A Rush to Medicate Young Minds. *Washington Post*, p. B07. Washington, D.C.
- Robillard, J. M., & Illes, J. (2011). Lost in translation: neuroscience and the public. *Nature Reviews Neuroscience*, 12, 118–118.
- Rosch, E., & Lloyd, B. B. (Eds.). (1978). *Cognition and categorization*. Hillsdale, N.J.: L. Erlbaum Associates ; distributed by Halsted Press.

- Rose, N. (2006). Disorders Without Borders? The Expanding Scope of Psychiatric Practice. *BioSocieties*, 1, 465–484.
- Rose, N., & Abi-Rached, J. (2013). *Neuro: the new brain sciences and the management of the mind*. Princeton, N.J.: Princeton University Press.
- Rosen, L. D., Whaling, K., Rab, S., Carrier, L. M., & Cheever, N. A. (2013). Is Facebook creating “iDisorders”? The link between clinical symptoms of psychiatric disorders and technology use, attitudes and anxiety. *Computers in Human Behavior*, 29, 1243–1254.
- Rosenblatt, F. (1958a). The perceptron: a probabilistic model for information storage and organization in the brain. *Psychological Review*, 65, 386.
- Rosenblatt, F. (1958b). *Two theorems of statistical separability in the perceptron (Project Para)*. Buffalo, New York: Cornell Aeronautical Laboratory, Inc.
- Rosenblatt, F. (1960). Perceptron Simulation Experiments. *Proceedings of the IRE*, 48, 301–309.
- Rosenblueth, A., Wiener, N., & Bigelow, J. (1943). Behavior, Purpose and Teleology. *Philosophy of Science*, 10, 18–24.
- Rumelhart, D. E., & McClelland, J. L. (1989). *Parallel Distributed Processing, Vol. 1: Foundations*. Cambridge, Mass.: MIT Press.
- Rushe, D. (2014, October 2). Facebook sorry – almost – for secret psychological experiment on users. Retrieved March 27, 2015, from <http://www.theguardian.com/technology/2014/oct/02/facebook-sorry-secret-psychological-experiment-users>
- Sahakian, B., & Morein-Zamir, S. (2007). Professor’s little helper. *Nature*, 450, 1157–1159.
- Sahlins, M. (1976). *The use and abuse of biology: an anthropological critique of sociobiology*. Ann Arbor: University of Michigan Press.
- Sample, I. (2014, July 6). Scientists threaten to boycott €1.2bn Human Brain Project. *The Guardian*. Retrieved from <http://www.theguardian.com/science/2014/jul/07/human-brain-project-researchers-threaten-boycott>
- Sanchez, J. (2015). Systems-Based Neurotechnology for Emerging Therapies (SUBNETS). Retrieved October 17, 2015, from <http://www.darpa.mil/program/systems-based-neurotechnology-for-emerging-therapies>
- Satel, S. (2003, July 23). ADD Overdose? *Wall Street Journal*. Retrieved from <http://online.wsj.com/article/SB105892212239632200.html>
- Satel, S. L. (2013). *Brainwashed: the seductive appeal of mindless neuroscience*. New York: Basic Books.
- Scassellati, C., Bonvicini, C., Faraone, S. V., & Gennarelli, M. (2012). Biomarkers and attention-deficit/hyperactivity disorder: a systematic review and meta-analyses. *Journal of the American Academy of Child and Adolescent Psychiatry*, 51, 1003–1019.e20.
- Schäfer, M. S. (2012). Taking stock: A meta-analysis of studies on the media’s coverage of science. *Public Understanding of Science*, 21, 650–663.
- Schaffer, S. (1997). Babbage’s Dancer and the Impresarios of Mechanism. In F. Spufford (Ed.), *Cultural Babbage: technology, time and invention* (pp. 53–80). London: Faber & Faber.
- Schrag, P., & Divoky, D. (1976). *The myth of the hyperactive child: and other means of child control* /. Dell.
- Schüll, N. D. (2012). Gambled Away: Video Poker and Self-Suspension. *Anthropology Now*, 4. Retrieved from <http://128.138.229.8/an/subscriptions/AnthroNowFall2012.pdf>
- Schwab-Stone, M. E., Shaffer, D., Dulcan, M. K., Jensen, P. S., Fisher, P., Bird, H. R., . . . Rae, D. S. (1996). Criterion Validity of the NIMH Diagnostic Interview Schedule for Children Version 2.3 (DISC-2.3). *Journal of the American Academy of Child & Adolescent Psychiatry*, 35, 878–888.

- Schwarz, A. (2012, October 9). Attention Disorder or Not, Children Prescribed Pills to Help in School. *The New York Times*. Retrieved from <http://www.nytimes.com/2012/10/09/health/attention-disorder-or-not-children-prescribed-pills-to-help-in-school.html>
- Schwarz, A. (2013a, February 2). Concerns About A.D.H.D. Practices and Amphetamine Addiction. *The New York Times*. Retrieved from <http://www.nytimes.com/2013/02/03/us/concerns-about-adhd-practices-and-amphetamine-addiction.html>
- Schwarz, A. (2013b, December 14). The Selling of Attention Deficit Disorder. *The New York Times*. Retrieved from <http://www.nytimes.com/2013/12/15/health/the-selling-of-attention-deficit-disorder.html>
- Schwarz, A. (2014, May 16). Thousands of Toddlers Are Medicated for A.D.H.D., Report Finds, Raising Worries. *The New York Times*. Retrieved from <http://www.nytimes.com/2014/05/17/us/among-experts-scrutiny-of-attention-disorder-diagnoses-in-2-and-3-year-olds.html>
- Sciar, D. A., Robison, L. M., Bowen, K. A., Schmidt, J. M., Castillo, L. V., & Oganov, A. M. (2012). Attention-Deficit/Hyperactivity Disorder Among Children and Adolescents in the United States Trend in Diagnosis and Use of Pharmacotherapy by Gender. *Clinical Pediatrics*, 51, 584–589.
- Searle, J. R. (1995). *The construction of social reality*. New York: Free Press.
- Sedgwick, P. (1981). Michel Foucault: The anti-history of psychiatry. *Psychological Medicine*, 11, 235–248.
- Seiss, D. (2011, June 19). Whatever Happened To....the boy with ADHD? *The Washington Post*. Retrieved from http://www.washingtonpost.com/lifestyle/magazine/whatever-happened-to-the-boy-with-adhd/2011/04/29/AGqWQ8XH_story.html
- Seung, S. (2012, June 11). Another Perspective on Massive Brain Simulations. Retrieved June 6, 2013, from <http://www.scientificamerican.com/article.cfm?id=massive-brain-simulators-seung-conntectome&page=2>
- Seung, S. (2013). *Connectome: how the brain's wiring makes us who we are* (First Mariner Books edition). Boston: Mariner Books, Houghton Mifflin Harcourt.
- Shaffer, D. (2000). NIMH Diagnostic Interview Schedule for Children Version IV (NIMH DISC-IV): Description, Differences From Previous Versions, and Reliability of Some Common Diagnoses. *Journal of the American Academy of Child & Adolescent Psychiatry*, 39, 28–38.
- Shannon, C. E. (1948). A Mathematical Theory of Communication. *Bell System Technical Journal*, 27, 379–423.
- Shapiro, M. (2007, January 2). How to recognize the hyper child. *Globe and Mail*, p. A11. Toronto.
- Shils, E. A., & Janowitz, M. (1948). Cohesion and Disintegration in the Wehrmacht in World War II. *Public Opinion Quarterly*, 12, 280–315.
- Shorter, E. (2008). *Before Prozac: The Troubled History of Mood Disorders in Psychiatry* (1st ed.). Oxford University Press, USA.
- Shrum, L. J. (2008). Media Consumption and Perceptions of Social Reality: Effects and underlying processes. In J. Bryant & D. Zillmann (Eds.), *Media Effects: Advances in Theory and Research* (pp. 69–95).
- Shrum, L. J., & O'guinn, T. C. (1993). Processes and Effects in the Construction of Social Reality Construct Accessibility as an Explanatory Variable. *Communication Research*, 20, 436–471.
- Simon, H. A. (1996a). *Models of my life* (1st MIT Press ed). Cambridge, Mass: MIT Press.
- Simon, H. A. (1996b). *The Sciences of the Artificial* (Third edition). Cambridge, Mass.: The MIT Press.

- Simon, H. A. (1997). *Administrative Behavior* (4th Revised edition). New York: Free Press.
- Simon, H. A., & Newell, A. (1958). Heuristic Problem Solving: The Next Advance in Operations Research. *Operations Research*, 6, 1–10.
- Singh, I. (2004). Doing their jobs: mothering with Ritalin in a culture of mother-blame. *Social Science & Medicine*, 59, 1193–1205.
- Slezak, P. (1991). How Strong Is the “Strong Programme”? *Social Studies of Science*, 21, 154–156.
- Slezak, P. (1992). Review of H. Collins, *Artificial Experts*. *Social Studies of Science*, 22, 175–201.
- Sloman, S. A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin*, 119, 3–22.
- Small, G. W., & Vorgan, G. (2009). *iBrain: surviving the technological alteration of the modern mind*. New York: Harper.
- Smith, M. (2011). *An Alternative History of Hyperactivity: Food Additives and the Feingold Diet*. Rutgers University Press.
- Smith, M. (2012). *Hyperactive: a history of ADHD*. London: Reaktion Books.
- Smolensky, P. (1988). On the proper treatment of connectionism. *Behavioral and Brain Sciences*, 11, 1–23.
- Snow, D. (2010, February 27). ADHD pill becomes students’ party, smart drug of choice. *Sydney Morning Herald*, p. 2. Sydney, Aus.
- Snyder, T. (2015). *Black Earth: the Holocaust as history and warning*. New York: Tim Duggan Books.
- Somers, J. (2013, November). The Man Who Would Teach Machines to Think. *The Atlantic*. Retrieved from <http://www.theatlantic.com/magazine/archive/2013/11/the-man-who-would-teach-machines-to-think/309529/>
- Sommers, C. H. (2001). *The WAR AGAINST BOYS: How Misguided Feminism Is Harming Our Young Men* (1st Touchstone Ed edition). New York: Simon & Schuster.
- Sparrow, B., Liu, J., & Wegner, D. M. (2011). Google Effects on Memory: Cognitive Consequences of Having Information at Our Fingertips. *Science*, 333, 776–778.
- Star, S. L. (2010). This is Not a Boundary Object: Reflections on the Origin of a Concept. *Science, Technology & Human Values*, 35, 601–617.
- Star, S. L., & Griesemer, J. R. (1989). Institutional Ecology, ‘Translations’ and Boundary Objects: Amateurs and Professionals in Berkeley’s Museum of Vertebrate Zoology, 1907–39. *Social Studies of Science*, 19, 387–420.
- Steers, M.-L. N., Wickham, R. E., & Acitelli, L. K. (2014). Seeing Everyone Else’s Highlight Reels: How Facebook Usage is Linked to Depressive Symptoms. *Journal of Social and Clinical Psychology*, 33, 701–731.
- Stiegler, B. (2010). *Taking care of youth and the generations*. Stanford, Calif: Stanford University Press.
- Still, G. F. (2006). Some Abnormal Psychical Conditions in Children: Excerpts from Three Lectures. *Journal of Attention Disorders*, 10, 126–136.
- Strathern, M. (1992). *Reproducing the Future: Anthropology, Kinship, and the New Reproductive Technologies*. New York: Routledge.
- Strauss, A. A., & Lehtinen, L. E. (1947). *Psychopathology and education of the brain-injured child*. Oxford, England: Grune & Stratton.
- Strum, S. S., & Latour, B. (1987). Redefining the social link: from baboons to humans. *Social Science Information*, 26, 783–802.
- Suchman, L. (1987). *Plans and Situated Actions: The Problem of Human-Machine Communication*. Cambridge: Cambridge University Press.
- Suchman, L. (1988). Representing practice in cognitive science. *Human Studies*, 11, 305–325.

- Suchman, L. (1993). Response to Vera and Simon's Situated Action: A Symbolic Interpretation. *Cognitive Science*, 17, 71–75.
- Sulzer, D., Chen, T. K., Lau, Y. Y., Kristensen, H., Rayport, S., & Ewing, A. (1995). Amphetamine redistributes dopamine from synaptic vesicles to the cytosol and promotes reverse transport. *The Journal of Neuroscience*, 15, 4102–4108.
- Sunstein, C. R. (2005). *Laws of Fear: Beyond the Precautionary Principle*. Cambridge University Press.
- Swing, E. L., Gentile, D. A., Anderson, C. A., & Walsh, D. A. (2010). Television and video game exposure and the development of attention problems. *Pediatrics*, 126, 214–221.
- Szalavitz, M. (2013, December 11). The boy whose brain could unlock autism. Retrieved October 23, 2014, from <https://medium.com/matter-archive/the-boy-whose-brain-could-unlock-autism-70c3d64ff221>
- Szasz, T. (2009). *Antipsychiatry quackery squared*. Syracuse, N.Y.: Syracuse University Press. Retrieved from <http://site.ebrary.com/id/10493602>
- Talbot, M. (2009, September 20). Can a daily pill really boost brain power? *The Observer*, p. 18. London, England.
- Thagard, P. (2012). *The cognitive science of science: explanation, discovery, and conceptual change*. Cambridge, Mass: MIT Press.
- The Neurocritic. (2013, October 29). A Tale of Two BRAINS: #BRAINI and DARPA's SUBNETS. Retrieved October 17, 2015, from <http://neurocritic.blogspot.ca/2013/10/a-tale-of-two-brains-braini-and-darpas.html>
- The word: Common sense. (2006, April 15). Retrieved October 6, 2014, from <http://www.newscientist.com/article/mg19025471.700-the-word-common-sense.html>
- Thome, J., & Jacobs, K. A. (2004). Attention deficit hyperactivity disorder (ADHD) in a 19th century children's book. *European Psychiatry: The Journal of the Association of European Psychiatrists*, 19, 303–306.
- Thompson, C. (2014, October 23). Can Video Games Fend Off Mental Decline? *The New York Times*. Retrieved from <http://www.nytimes.com/2014/10/26/magazine/can-video-games-fend-off-mental-decline.html>
- Timimi, S., & Taylor, E. (2004). ADHD is best understood as a cultural construct. *The British Journal of Psychiatry*, 184, 8–9.
- Turing, A. (1936). On computable numbers, with an application to the Entscheidungsproblem. *J. of Math*, 58, 345–363.
- Turing, A. (1950). Computing Machinery and Intelligence. *Mind*, 59, 433–460.
- Turkle, S. (Ed.). (2007). *Evocative objects: things we think with*. Cambridge, Mass: MIT Press.
- Turn off the television. (2004, April 11). *The Japan Times*. Tokyo, Jpn.
- TV “rewires” developing brains, researchers fear. (2004, April 5). *Washington Times*.
- Ulrich, M., Keller, J., Hoenig, K., Waller, C., & Grön, G. (2014). Neural correlates of experimentally induced flow experiences. *NeuroImage*, 86, 194–202.
- Uribe, R., & Gunter, B. (2004). Research Note: The Tabloidization of British Tabloids. *European Journal of Communication*, 19, 387–402.
- Uttal, W. R. (2003). *The New Phrenology: The Limits of Localizing Cognitive Processes in the Brain*. MIT Press.
- Uttal, W. R. (2011). *Mind and brain: a critical appraisal of cognitive neuroscience*. Cambridge, Mass: MIT Press.
- Vallée, M. (2002). Resisting American Psychiatry: French Opposition to DSM-III, Biological Reductionism, and the Pharmaceutical Ethos. *Advances in Medical Sociology*, 12, 85–110.

- Varela, F. J., Thompson, E., & Rosch, E. (1991). *The embodied mind: cognitive science and human experience*. Cambridge, Mass.: MIT Press.
- Vera, A. H., & Simon, H. A. (1993). Situated Action: A Symbolic Interpretation. *Cognitive Science*, 17, 7–48.
- Vidal, F. (2009). Brainhood, anthropological figure of modernity. *History of the Human Sciences*, 22, 5–36.
- Vincent, A., & Hoff, H. S. (2010). *My brain needs glasses: living with hyperactivity*. Montréal: Éditions Québecor.
- Von Neumann, J. (1956). Probabilistic logics and the synthesis of reliable organisms from unreliable components. *Automata Studies*, 34, 43–98.
- Von Neumann, J. (1966). *Theory of self-reproducing automata*. (A. W. Burks, Ed.). Urbana-Champaign: University of Illinois Press. Retrieved from <http://books.google.ca/books?id=ktNQAAAAMAAJ>
- Von Neumann, J. (2012). *The computer & the brain* (3rd ed). New Haven, Conn.: Yale University Press.
- Vul, E., Harris, C., Winkielman, P., & Pashler, H. (2009). Puzzlingly High Correlations in fMRI Studies of Emotion, Personality, and Social Cognition. *Perspectives on Psychological Science*, 4, 274–290.
- Waldrop, M. M. (2001). *The dream machine: J. C. R. Licklider and the revolution that made computing personal*. New York: Viking.
- Waldrop, M. M. (2012). Computer modelling: Brain in a box. *Nature*, 482, 456–458.
- Walter, W. G. (1953). *The living brain*. New York: W. W. Norton.
- Walton, D. (2001, April 7). Kiddie coke; Last year Canadian doctors prescribed almost 78 million doses of the potent stimulants used to combat attention-deficit disorder. *Globe and Mail*, p. F4. Toronto, Can.
- Ware, W. (1953, March 10). The history and development of the electronic computer project at the Institute for Advanced Study. RAND Corporation.
- Watters, E. (2011). *Crazy Like Us: The Globalization of the American Psyche* (Reprint). Free Press.
- Webb, L. J. (Ed.). (1981). *DSM-III training guide for use with the American Psychiatric Association's Diagnostic and statistical manual of mental disorders (third edition)*. New York: Brunner/Mazel.
- Weber, R., Ritterfeld, U., & Mathiak, K. (2006). Does Playing Violent Video Games Induce Aggression? Empirical Evidence of a Functional Magnetic Resonance Imaging Study. *Media Psychology*, 8, 39–60.
- Weber, R., Tamborini, R., Westcott-Baker, A., & Kantor, B. (2009). Theorizing Flow and Media Enjoyment as Cognitive Synchronization of Attentional and Reward Networks. *Communication Theory*, 19, 397–422.
- Wedge, M. (2012, March 8). Why French Kids Don't Have ADHD. Retrieved November 22, 2013, from <http://www.psychologytoday.com/blog/suffer-the-children/201203/why-french-kids-dont-have-adhd>
- Weisberg, D. S., Keil, F. C., Goodstein, J., Rawson, E., & Gray, J. R. (2007). The Seductive Allure of Neuroscience Explanations. *Journal of Cognitive Neuroscience*, 20, 470–477.
- Wheeler, J. A. (1990). "It from Bit" as a Guide in the Search for Links Connecting Physics, Quantum, and Information. In W. H. Zurek (Ed.), *Complexity, Entropy, and the Physics of Information: The Proceedings of the 1988 Workshop on Complexity, Entropy, and the Physics of Information*. Boulder, Co.: Westview Press.

- White, B. P., Becker-Blease, K. A., & Grace-Bishop, K. (2006). Stimulant Medication Use, Misuse, and Abuse in an Undergraduate and Graduate Student Sample. *Journal of American College Health, 54*, 261–268.
- White, N. J. (2008, May 20). Adult ADHD tough to diagnose. *The Toronto Star*. Retrieved from http://www.thestar.com/life/health_wellness/diseases_cures/2008/05/20/adult_adhd_tough_to_diagnose.html
- Wiener, J. B., & Rogers, M. D. (2002). Comparing precaution in the United States and Europe. *Journal of Risk Research, 5*, 317–349.
- Wiener, N. (1961). *Cybernetics, or, Control and communication in the animal and the machine* (2nd ed.). Cambridge, Mass.: M.I.T. Press.
- Wiener, N. (1964). *God and Golem, Inc.: A commentary on certain points where cybernetics impinges on religion*. Cambridge, Mass.: MIT Press.
- Wiener, N. (1967). *The human use of human beings; cybernetics and society*. London: Avon Books.
- Wiener, N. (1993). *Invention: the care and feeding of ideas*. Cambridge, Mass: MIT Press.
- Wilms, M., Schilbach, L., Pfeiffer, U., Bente, G., Fink, G. R., & Voegeley, K. (2010). It's in your eyes—using gaze-contingent stimuli to create truly interactive paradigms for social cognitive and affective neuroscience. *Social Cognitive and Affective Neuroscience*, nsq024.
- Wilson, A. D., & Golonka, S. (2013). Embodied Cognition is Not What you Think it is. *Frontiers in Psychology, 4*. doi:10.3389/fpsyg.2013.00058
- Winograd, T. (1971). *Procedures as a Representation for Data in a Computer Program for Understanding Natural Language (manuscript)*. Cambridge, Mass.
- Winograd, T. (1991). Oral history interview with Terry Allen Winograd. Retrieved from <http://conservancy.umn.edu/handle/11299/107717>
- Winograd, T., & Flores, F. (1986). *Understanding Computers and Cognition: A New Foundation for Design*. Bristol: Intellect Books.
- Womack, S. (2006, April 21). Spare the “chemical cosh”; Baroness warns of a growing scandal over the number of youngsters prescribed Ritalin. *Daily Telegraph*, p. 8. London.
- Yeo, R. R. (2003). *Defining science: William Whewell, natural knowledge and public debate in early Victorian Britain*. Cambridge: Cambridge University Press.
- Yong, E., Ledford, H., & Van Noorden, R. (2013). Research ethics: 3 ways to blow the whistle. *Nature, 503*, 454–457.

Appendix: Figure references and permissions

- Figure 1:* Keith (1920, p. 258). Public domain image, over 50 years elapsed since death of author.
- Figure 2:* Image from Wikimedia Commons, reproduced under Creative Commons Attribution-Share Alike license. https://commons.wikimedia.org/wiki/File:Cognitive_Science_Hexagon.svg – derivative work, *cf.* (Miller, 2003, p.143).
- Figure 3:* Manipulated digital photo-collage by Suzanne Treister, part of artistic installation *Hexen 2.0*. <http://www.suzannetreister.net/HEXEN2/Seance/cyberneticseance.html>. Reproduced with permission of the artist.
- Figure 4:* W. Ross Ashby's journals, p.2431. Reproduced under blanket permission for scholarly non-commercial use, <http://rossashby.info/copyright.html>
- Figure 5:* Image from W. Ross Ashby online archive. Reproduced under blanket permission for scholarly non-commercial use, <http://rossashby.info/copyright.html>
- Figure 6:* Latil (1956, p. 297). Orphan work, original publisher Houghton Mifflin advises they no longer hold copyright; they directed me to Elsevier, who advised in turn that they do not hold copyright and referred me back to Houghton Mifflin. Date of author's death unknown. Reproduced under fair dealing exemption for the purposes of scholarly review/comment.
- Figure 7:* Latil (1956, p. 299). See note for *fig. 6* above.
- Figure 8:* Latil (1956, p. 273). See note for *fig. 6* above.
- Figure 9:* W. Ross Ashby's own photo, from online archive. Reproduced under blanket permission for scholarly non-commercial use, <http://rossashby.info/copyright.html>
- Figure 10:* Latil (1956, p. 275, p. 34). See note for *fig. 6* above.
- Figure 11:* Pryor (1961, p. 19). Declassified U.S. Government document, public domain.
- Figure 12:* Hay, Lynch, Smith, & Murray (1960, p. 4). Declassified U.S. Government document, public domain.
- Figure 13:* Cover of Minsky and Papert (1969). © 1969, MIT Press, reproduced with permission of copyright holder.
- Figure 14:* Public domain image reproduced from Wikimedia Commons, https://commons.wikimedia.org/wiki/File:Arpanet_logical_map,_march_1977.png
- Figure 15:* Image from Herbert Simon papers, online archive, Carnegie Mellon University. Reproduced with permission of the CMU Archivist.
- Figure 16:* Newell and Shaw (1957, p. 234). © 1957, IEEE. Reproduced with permission of current copyright holder, IEEE, original rightsholder AIEE merged into this organization.
- Figure 17:* Newell, Shaw & Simon(1958b, p. 21). Reproduced with permission of RAND Corporation.
- Figure 18:* Winograd (1971, p. 25). © 1971, Terry Winograd, dissertation manuscript. Reproduced with permission of the author.
- Figure 19:* Klatt (1977, p. 1352). © 1977, reproduced with permission of the Acoustical Society of America.

- Figure 20:* Hofstadter and Mitchell, in Hofstadter, ed. (1996, p. 238). © 1996 Basic Books, reproduced with permission of copyright holder.
- Figure 21:* Brooks (1991, p. 152). © 1991, Elsevier, reproduced with permission of copyright holder.
- Figure 22:* Markram (2012, p. 37). © 2012, Scientific American; Scientific American grants non-exclusive license for reuse of materials for educational purposes.
<http://www.scientificamerican.com/pressroom/reprints-permissions/>. Graph in this figure reproduces one from (Waldrop, 2012); permission to reproduce also granted by original copyright holder, © 2012, Springer Nature.
- Figure 23:* Diagram from Human Brain Project press kit, reproduced under blanket permission for republication of these materials: <https://www.humanbrainproject.eu/platform-release>
- Figure 24:* Infographic image from IBM Research Web site,
http://research.ibm.com/software/IBMResearch/image/brain_banner_infographic.jpg. © 2015, IBM, reproduced with permission of the copyright holder.
- Figure 25:* Eliasmith *et al.*, (2012, p. 1204). © 2012, AAAS, reproduced with permission; blanket permission for reuse of images in theses/dissertations:
<http://www.sciencemag.org/help/reprints-and-permissions>
- Figure 26:* Gerbner *et al.*, (1979, p. 195). © 1979, John Wiley and Sons, reproduced with permission of copyright holder.
- Figure 27:* Mathiak and Weber (2006, pp. 952, 953). © 2006, John Wiley and Sons, reproduced with permission of copyright holder.
- Figure 28:* Christakis *et al.*, (2012, p. 2). Reproduced under Creative Commons (Attribution-Noncommercial) license.
- Figure 29:* From Lumosity press materials, <https://www.lumosity.com/press/resources>. Reproduced with permission.
- Figure 30:* Gameplay screenshot reproduced under Creative Commons license from Wikimedia Commons: https://commons.wikimedia.org/wiki/File:Beckee_Morrison_626_points!.jpg
- Figure 31:* Artist's rendering produced by DARPA contractor, retrieved from
<http://www.doncio.navy.mil/uploads/0527MHO97882.jpg>. Reproduced under blanket permission for educational/informational use, <http://www.darpa.mil/policy/usage-policy>
- Figure 32:* Bradley and Bowen (1941, pp. 95, 99). © 1941, American Psychological Association, reproduced under blanket permission for noncommercial use of less than three figures,
<http://www.apa.org/about/contact/copyright/index.aspx>
- Figure 37:* Graph reproduced under Creative Commons license from Wikimedia Commons:
<https://commons.wikimedia.org/wiki/File:Lorenz.png>