MEDIAL FRONTAL CORTEX NEURONAL ACTIVITIES DURING REPEATED RULE SWITCHES IN RATS

RYLEY NATHANIEL

A THESIS SUBMITTED TO THE FACULTY OF GRADUATE STUDIES

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

FOR THE DEGREE OF

MASTER OF ARTS

GRADUATION PROGRAM IN PSYCHOLOGY

YORK UNIVERSITY

TORONTO, ONTARIO

JULY 2026

© RYLEY NATHANIEL, 2026

Abstract

Rule-switching tasks are routinely used to study the neural underpinnings of cognitive flexibility. However, existing paradigms often lack the trial density or temporal precision—specifically a neutral delay period—required to isolate neural correlates of rule-guided decision-making. In this study, we bridge this gap while introducing a novel feedback-driven Rule-Switching task in which rats learned to switch between blocks of trials under one of two rules. In each trial, prior to the response period, rats nose-fixated in the central port for a 1-second delay period, during which the rule and previous choice(s) must be maintained. We demonstrated that as rats transition from Beginner to Expert over 16 days of training, their behavioural improvements are not mirrored by "sharper" individual neuronal signals. Instead, behavioural mastery is underpinned by an expansion in population-level dynamic neural recruitment, specifically within ACC-M2 areas. Furthermore, we found that Expert animals committed significantly fewer perseverative errors. This work provides a new framework for understanding cognitive control as a distributed persistent population state rather than a static individual neuron property.

Keywords: cognitive flexibility, rule-switching, recruitment, density, intensity, mastery, win-Stay, win-Shift, perseverative error, regressive error

Table of Contents

Abstract ii
Table of Contents iii
Table of Equations v
Table of Figures vi
Table of Tables vii
1 General background 1
1.1 Paradigms assessing cognitive flexibility and flexibility impairments 3
1.2 Methodological gap in assessing cognitive flexibility 10
1.3 Addressing the neurophysiological knowledge gap 12
1.4 Our current task & goal 14
2 Rats perform multiple switches on rule-switching task 16
2.1 Background 16
2.2 Methods 18
2.2.1 Animals 18
2.2.2 Surgery 18
2.2.3 Apparatus 19
2.2.4 Histology 20
2.2.5 Task training 22
2.2.6 Data analyses 24
2.3 Results 27
2.4 Discussion 35
2.4.1 The evolution of the rat error landscape 36
2.4.2 Information-theoretic dissociation of performance metrics 37
3 Learning-induced population recruitment in ACC-M2 networks 38
3.1 Background 38
3.2 Methods 41
3.2.1 Spike sorting 41
3.2.2 Data analyses 41
3.3 Results 44
3.4 Discussion 58

3.4.1 Overlapping latent spaces and sensorimotor gating 59
3.4.2 Population expansion and dynamic neural recruitment over metabolic extremes 59
4 General Discussion 62
4.1 Statistical predictors of task success: Trial Epoch, Rule type, and Stage 65
4.2 Reaching the cognitive ceiling: Reaction time and practice 67
4.3 From perseveration to anticipation: Error analysis 68
4.4 Rapid communication: Similar population coding of task variables 70
4.5 The Relay Race of information: Dynamic neural recruitment as a learning strategy 72
4.6 A unified frontal circuit: Asymmetric network optimization 73
4.7 Limitations 74
4.8 Relevance and application 77
4.9 Future directions 78
4.10 Conclusion 81

Table of Equations

(1) 26

(2) 26

(3) 26

(4) 44

(5) 44

Table of Figures

Figure 1 20

Figure 2 24

Figure 3 29

Figure 4 31

Figure 5 33

Figure 6 34

Figure 7 35

Figure 8 47

Figure 9 49

Figure 10 53

Figure 11 57

Table of Tables

Table 1 28

Table 2 30

Chapter 1. General background

Throughout daily life, humans are faced with countless environmental and/or contextual changes, and we must flexibly select an appropriate response in each situation. For example, when traffic lights turn yellow, drivers must decide whether the remaining distance to the intersection is sufficient for braking. Perseverative adoption of the same response (i.e., continuing to press the gas pedal) without reassessment of the situation could have detrimental consequences. This ability to adapt behaviour in response to changing task demands is a critical component of cognitive control called *cognitive flexibility*. Cognitive flexibility is the foundation of adaptive behaviour, allowing us to pivot our strategies when the environment or context changes. When it comes to researching cognitive flexibility, mapping its high-level human adaptations onto precise neural mechanisms presents a major methodological and translational challenge. To resolve how the brain coordinates these environmental/contextual changes, neuroscience must balance the structural complexity of a behavioural model with the capacity for longitudinal circuit recordings.

As the group of mammals closest to humans, apart from non-human primates (NHPs), rodents—especially rats—have proven to be a suitable model for understanding the neural basis of cognitive control. The rat brain contains anterior cingulate areas that are homologous to ours. Additionally, while they do not have a frontal eye field (FEF), the neck area of the secondary motor cortex (M2) is believed to contain a frontal orienting field (FOF), which may contribute to spatial attention and planning (Barthas & Kwan, 2017; Donoghue & Wise, 1982; Erlich, Bialek, & Brody, 2011; Erlich, Brunton, Duan, Hanks, & Brody, 2015). The rat anterior cingulate cortex

(ACC) and M2 effectively mirror the functional sub-components of the primate executive system (Laubach, Amarante, Swanson, & White, 2018; Uylings, Groenewegen, & Kolb, 2003). Overall, this allows us to cleanly capture the precise cellular transformation that occurs as an animal dynamically updates its behaviour strategy during cognitive flexibility paradigms.

Another consideration when researching cognitive flexibility is its presentation of a sophisticated computational challenge for the brain, requiring tight integration between cognitive evaluation and motor execution. This challenge is particularly pronounced during adaptive behaviour, where abstract internal states (e.g., reward history, error monitoring) must be rapidly translated into precise behavioural outputs. While the ACC is traditionally viewed as an evaluator of task effort, conflict, and outcome (Kolling, Behrens, Wittmann, & Rushworth, 2016) and the M2 as a center for action preparation (Erlich, Bialek, & Brody, 2011), the dynamic interaction between these regions during flexible learning remains elusive. Historically, classic architectures such as the *Conflict-Monitoring Model* (Botvinick, Braver, Barch, Carter, & Cohen, 2001) have conceptualized these components as a serial (i.e., linear) processing pipeline, where the ACC's evaluative signals must be finalized before being passed downstream to update M2 selection rules. However, real-time adaptation demands continuous bidirectional communication between the regions.

Since keeping up constant communication between the ACC and M2 regions requires a lot of energy, flexible learning is often modelled based on the *Efficient Coding Hypothesis* framework (previously the *Redundancy Reduction Hypothesis*), which suggests that expertise should sharpen neural tuning to reduce metabolic cost (Barlow, 1961). Conversely, newer frameworks suggest that robustness in complex environments requires dynamic neural recruitment (Friedman & Robbins, 2022; Miller & Cohen, 2001). In the present study, we

evaluate these competing hypotheses by leveraging high-density rat recordings to track the ACC-M2 circuit dynamics, specifically during the inter-trial learning window (i.e., the critical bridge between past feedback and future preparation). By capturing these transient network reorganizations during our novel Rule-Switching task, we address a critical gap in current models that typically overlook this vital window of strategy acquisition.

1.1 Paradigms assessing cognitive flexibility and flexibility impairments

The Wisconsin Card Sorting Test (WCST): Mechanics, circuits, and limitations

In neuropsychology, cognitive flexibility is classically assessed using paradigms that require the subject to map changing environmental feedback to internal goal states. The most widely used attentional set-shifting task for measuring this flexibility in humans is the Wisconsin Card Sorting Test, or WCST (Grant & Berg, 1948). The WCST is designed to measure rule-switching abilities by requiring participants to adapt their response strategies through trial-and-error. Participants are presented with reference and test cards that they must sort according to one of three hidden dimensions: colour, shape, or number (Konishi, et al., 1999; Monchi, Petrides, Petre, Worsley, & Dagher, 2001). This sorting rule alternates intermittently without warning and is never explicitly revealed; instead, participants must deduce the active rule entirely from feedback indicating whether each match is correct.

Mastery of the WCST and its complex dimensional transitions requires the ACC to detect error signals when a previously correct rule leads to failure, and the prefrontal cortex (PFC) to switch the mental set to a new dimension (Botvinick, Braver, Barch, Carter, & Cohen, 2001; Carter, et al., 1998). Meanwhile, impaired card sorting is seen in patients with lesions of the PFC

(Milner, 1963; Nelson, 1976; Stuss, et al., 2000). Impaired performance on the WCST has further been linked to various clinical conditions, including schizophrenia (Rady, Elsheshai, el Wafa, & Elkholy, 2012; Wang, Hu, & Li, 2022). Crucially, specialized cognitive set-shifting deficits observed during these paradigms are directly tied to phenotypic symptom profiles; for instance, Miller and colleagues (2015) demonstrated that an inability to maintain and selectively update response sets in clinical groups tracks significantly with the severity of real-world restricted and repetitive behaviour. Poor cognitive flexibility has been reported in additional neuropsychiatric conditions such as eating disorders (Cochereau, Zandvoort, Santariius, & Mandonnet, 2021; Milton, et al., 2021; Perpiñá, Segura, & Sánchez-Reales, 2017; Zastrow, et al., 2009), major depressive disorder (Murphy, Michael, & Sahakian, 2012; Murphy, et al., 1999), and anxiety disorders (Park & Moghaddam, 2017; Yu, Yu, & Lin, 2020).

Furthermore, primate lesion and pharmacological inactivation studies have established a definitive causal link between specific prefrontal nodes and rule-switching capabilities. Reversible chemical inactivation of the macaque inferior arcuate region via focal muscimol microinjections selectively disrupts behavioural flexibility (Kamigaki, Fukushima, Tamura, & Miyashita, 2012). The severity of the shifting deficit tightly tracks the pharmaceutical dose (Kamigaki, Fukushima, Tamura, & Miyashita, 2012). Similarly, selective bilateral lesions within distinct prefrontal subdivisions reveal anatomical divisions of labour: while a broad distribution of prefrontal areas does not govern default dimensional biases, targeted lesions within the ACC, dorsolateral PFC, or orbitofrontal cortex (OFC) severely impair a primate's capacity to execute structural shifts between visual dimensions (Mansouri, Buckley, Fehring, & Tanaka, 2020).

To translate the WCST paradigm into animal models, automated analogs like the CANTAB Intra-dimensional Shift (IDS)/Extra-dimensional Shift (EDS) task (Cambridge Cognition, Ltd),

are widely used. These tasks demand a higher order of cognitive abstraction when compared to more basic behavioural tests like simple reversal learning (Reindl, et al., 2023). In these environments, stimuli possess multiple, overlapping feature dimensions (e.g., colour, shape, line orientation, or spatial location), forcing the subject to formulate an internal attentional set to prioritize one dimension (e.g., shape) while ignoring irrelevant distractors (e.g., colour). During an IDS (Dias, Robbins, & Roberts, 1996), new exemplars are introduced, but the rule remains restricted to the same broad dimension (e.g., matching a novel shape while ignoring colour). Conversely, during an EDS (Dias, Robbins, & Roberts, 1996; Owen, Roberts, Polkey, Sahakian, & Robbins, 1991), the overarching rule undergoes a complete structural update, transitioning from one entire dimension to another (e.g., shifting from shape to colour). This requires a major cognitive leap, forcing the subject to shift attention to an entirely new, previously irrelevant sensory dimension. Successfully navigating an EDS means the NHP must fully deconstruct their prior mental framework, disengage attention from a historically potent stimulus category, and selectively re-allocate attentional resources to a previously ignored dimension.

Neurophysiological recordings during these dimensional transitions in NHPs reveal highly localized neural substrates dedicated to this process. For instance, when macaques are required to shift between colour-matching and shape-matching rules, a distinct population of shift-related neurons along the posterior bank of the inferior arcuate sulcus exhibits transient, selective activation specifically during the execution of the cognitive set shift (Kamigaki, Fukushima, Tamura, & Miyashita, 2012). Direct cross-species behavioural comparisons between humans and macaques executing identical set-shifting paradigms demonstrate that while both species possess an inherent baseline bias toward specific visual dimensions (e.g., humans biased toward colour, macaques toward shape), the ability to flexibly override these default biases and shift between

dimensions relies strictly on top-down prefrontal executive control (Mansouri, Buckley, Fehring, & Tanaka, 2020). While the baseline representation of competing visual features likely emerges from bottom-up sensory architectures, the executive capacity to disengage from one dimension and prioritize another is mediated by ACC-M2 networks.

This process also extends beyond the borders of the frontal lobe, recruiting a highly coordinated frontoparietal network. Recording single-unit activity from the posterior parietal cortex (PPC) during a macaque analog of the WCST revealed that shift-related networks exhibit transient, directional shift-related activation during rule reconfigurations (Kamigaki, Fukushima, & Miyashita, 2009). Crucially, this parietal shift-related activity emerges roughly four seconds before the physical execution of the behavioural response and accurately predicts whether a set shift will be completed successfully. Deconstructing this parietal signal further reveals that behavioural shifting is composed of dual, parallel processes: a preparatory phase as the animal is recognizing a rule change, and an execution phase when they update their choice. Rather than dividing these tasks across anatomically segregated structures, the PPC utilizes a functionally intermingled circuitry where most of single units encode both preparation and execution signals dynamically (Kamigaki, Fukushima, & Miyashita, 2011). This is the case whether driven by external feedback or internal spontaneous strategy changes. Thus, the primate literature frames this shifting as a distributed, highly dynamic network phenomenon across the ACC-M2 axis.

Along with its strengths, the WCST also has its inherent methodological limitations as a standalone measure of cognitive flexibility. As critically evaluated by Miles and colleagues (2021), the WCST bundles pure set-shifting ability with heavy, competing cognitive demands. This includes working memory loads, abstract concept formation, and acute feedback sensitivity. When a participant commits an error on this paradigm, it is challenging to parse whether the

failure stems from a true deficit in executive flexibility or simply from secondary confounds (Miles, et al., 2021). Confounds include, but are not limited to, transient lapses in working memory updating or adverse emotional reactions to negative feedback. These confounds mask the specific cognitive relationships under investigation, creating a need for targeted behavioural designs that isolate discrete executive components.

Task-switching, reversal, and dimensional shift as alternatives

Beyond the WCST, the Task-Switching Paradigm (Rogers & Monsell, 1995) provides a higher-resolution look at the switch cost—the measurable increase in reaction time and error rate that occurs when a rule changes. Rogers and Monsell (1995) designed this paradigm to isolate the cognitive mechanisms of executive control, by using explicitly defined instructions and predictable task transitions to quantify the precise mental cost of shifting cognitive sets. Their core innovation is the AABB pattern, where participants are exposed to an alternating sequence of two distinct tasks (Task A and Task B) performed on highly structured, ambiguous stimuli. Despite their distinct methodological designs, these two human paradigms share a fundamental cognitive architecture, both requiring the rapid suppression of a previously reinforced rule and the active retrieval of an alternate behavioural schema. Focusing on these shared mechanisms, these human paradigms collectively highlight a critical temporal dynamic: rule-switching-induced decision-making occurs during a delay period between the feedback of the previous trial and the presentation of the next stimulus. Rather than a passive transition, it is during this window that the human brain must dynamically bridge the gap between retrospective evaluation (i.e., what occurred on the last trial) and prospective configuration (i.e., what rules to apply next).

To map the precise neural circuitry driving this internal bridge, neuroscientists rely on comparative animal models that replicate these exact rule-switching dynamics. The most popular

rodent rule-switch tasks include: (1) the tactile-olfactory set-shifting task (tactile-olfactory task), a rodent IDS/EDS task where rodent's dig for buried food reward based on scent or texture (Birrell & Brown, 2000; Cho, et al., 2020); (2) the automated, operant box-based set-shifting task (operant-box task), in which rodents switch between a light-cue rule and a correct-side rule, depicting a win-stay (Floresco, Block, & Tse, 2008); and (3) maze-based tasks, in which rodents must navigate a physical track and dynamically shift their behavioural strategies mid-session between a spatial rule and a visual-cue rule (Ding, Tomsick, Young, & Jadhav, 2025; Ghods-Sharifi, Haluk, & Floresco, 2008; Rich & Shapiro, 2007; Rich & Shapiro, 2009). All these tasks take advantage of the sensory and cognitive strengths of rats and mice, requiring only 2-3 weeks of training.

On a macrostructural scale, many paradigms have established that distinct ACC-M2 subregions dictate specific facets of behavioural updating (Birrell & Brown, 2000; Cho, et al., 2020; Ding, Tomsick, Young, & Jadhav, 2025; Floresco, Block, & Tse, 2008; Ghods-Sharifi, Haluk, & Floresco, 2008; Rich & Shapiro, 2007; Rich & Shapiro, 2009; St. Onge & Floresco, 2010). Previous research has substantiated that the medial frontocortical networks govern strategy-level shifts, while orbitofrontal networks selectively track value updates and environmental feedback (Hamilton & Brigman, 2016). These structural properties follow a clear evolutionary trajectory, offering a translational baseline where comparative frameworks like the reversal learning index can be utilized to evaluate systemic rule adjustments across species (Rajalakshmi & Jeeves, 1965).

However, these tasks have their limitations, starting with the number of trials produced. The operant-box task is exempt from this limitation, typically producing 100-200 trials within a given session. In comparison, the labor-intensive nature of the tactile-olfactory task and maze-

based tasks results in them only producing up to dozens of trials, which is too few for single unit and ensemble analysis. When encompassing all three aforementioned tasks, we note the limitation that none of them had a movement- and posture-neutral delay period. A recently developed pro- and anti-orienting task (Duan, Erlich, & Brody, 2015; Duan, et al., 2021) incorporated such a delay period to disentangle neural correlates for rule-guided decision-making from the preparation and execution of the response itself. However, this task is distinct from the other rule-switch tasks, which are feedback-based. Another limitation is the presence of only one switch in the rule before animals lost motivation. The reduction in rule-switches bottlenecks the data, compromising statistical power, and making it more difficult to isolate true executive control from other confounding variables. Thus, cognitive flexibility can be assessed at multiple distinct levels: attentional set-shifting, reversal learning, IDS, and EDS (Brigman, Graybeal, & Holmes, 2010).

In our paradigm, we focus specifically on a different level: rule-switching. Rule-switching is defined here as the ability to alternate between abstract instructional contingencies without changing the underlying sensory features or the stimuli themselves. While lower-level shifts like reversal learning or IDS modify choices based on stimulus values or features, rule-switching represents a distinct cognitive demand with changes to the goal or instruction. This relies on the lateral PFC in primates (Dias, Robbins, & Roberts, 1996; Dias, Robbins, & Roberts, 1997; Owen, Roberts, Hodges, & Robbins, 1993; Owen, Roberts, Polkey, Sahakian, & Robbins, 1991) and the ACC in rodents (Birrell & Brown, 2000; Floresco, Block, & Tse, 2008). Additionally, rule-switching likely engages dynamics across an extended brain network, as per functional MRI studies on set-shifting and other human task-switching (Kramer, et al., 2008; Nakahara, Hayashi, Konishi, & Miyashita, 2002; Vallesi, Arbula, Capizzi, Causin, & D'Avella, 2015). How such

network-wide dynamics play out during a flexible switch in an abstract rule remains unknown. Answering this question will require simultaneous recordings from multiple brain regions in the network. Current tasks are not ideally suited for the analysis of the role of individual neurons in rule switches, since rats could not easily accomplish more than one switch within a continuously recorded session (Brady & Floresco, 2015). It is therefore necessary to develop a task in which rats can be trained with relative ease to achieve multiple rule switches.

1.2 Methodological gap in assessing cognitive flexibility

In NHP neurophysiology, cognitive flexibility has historically been interrogated using sophisticated behavioural paradigms (Cantwell, Buckholtz, Atencia, & Rosati, 2022). These paradigms can be broadly dissociated into two cognitive operations: (1) simple reinforcement-driven reversal learning and (2) abstract attentional set-shifting. The most widespread baseline assessment of flexibility in NHPs is the Deterministic Reversal Learning Task (Reindl, et al., 2023). In a typical configuration, a macaque (*Macaca mulatta*) or common marmoset (*Callithrix jacchus*) is presented with two or more distinct visual or spatial stimuli where one is predictive of a reward (e.g., juice) and the other remains unrewarded. Once the subject achieves an automated baseline of accurate choice execution, the underlying reward contingencies are abruptly inverted without warning or cue. To maximize reward intake, the animal must suppress the previously reinforced stimulus-response association and adaptively switch to the newly rewarded option. Variations of this paradigm introduce probabilistic feedback, where the optimal stimulus delivers an adverse outcome on a minor percentage of trials (e.g., an 80/20 schedule), forcing the NHP to distinguish between stochastic environmental noise and true structural change in the rules (Cantwell, Buckholtz, Atencia, & Rosati, 2022).

Despite literature characterizing flexibility in NHPs, a fundamental methodological blind spot remains. The primary limitation of traditional NHP neurophysiology is its reliance on overtrained subjects. Since macaques are exceptionally expensive to maintain and require months of intensive baseline shaping to interact with complex digital or physical apparatuses, neurophysiological recordings are almost universally conducted in animals that have already achieved near-ceiling task mastery. Consequently, NHP literature is biased towards capturing the stable, finalized neural correlates of automated task execution. The initial transition from Beginner behaviour to expertise and the structural shifting of neural manifolds during this developmental window remain largely hidden in traditional Old World primate models. Moreover, findings by Kamigaki et al. (2012) and Mansouri et al. (2020) underscore the ACC's critical causal role in behavioural flexibility, yet they highlight a pressing knowledge gap: so far there has only been reversible inactivation or lesion studies, thus neurophysiological recordings are needed to investigate how dynamics in an intact network reorganize to support flexibility during the learning process itself. This is specifically the case for cooperative populations across the ACC and downstream motor-preparatory regions like the M2.

Meanwhile, rats offer unparalleled precision for recording from large-scale population dynamics across weeks of training using high-density probes (e.g., Neuropixels) or multi-photon calcium imaging. Despite their prefrontal architecture lacking directly homologous regions found in NHPs, they present a highly tractably structured network of ACC-M2 areas that effectively mirror the functional, if not macro-anatomical, sub-components of the primate executive system.

The rat circuit choice is particularly well-suited to address the temporal and functional dynamics highlighted by primate frontoparietal studies. A study demonstrated that single-cell networks must fluidly multiplex abstract rule preparation and motor-preparatory execution

during a behavioural shift (Kamigaki, Fukushima, & Miyashita, 2011). In the rat brain, this functional division of labour is cleanly distributed across the ACC-M2 axis. Where the ACC operates as a primary evaluator of task rules and feedback outcomes, the M2 serves as a critical hub for action preparation and execution pipelines. Given that primate PPC units dynamically blend preparation and execution signals seconds before an action takes place (Kamigaki, Fukushima, & Miyashita, 2011), tracking the real-time population interactions between the ACC and M2 across learning blocks provides a circuit-level window into this process. It allows us to capture exactly how evaluative rule representations are mapped onto preparatory motor pipelines as an animal transitions from Beginner performance to baseline task mastery.

1.3 Addressing the neurophysiological knowledge gap

As highlighted above, neurophysiological studies in animals provide insight to how key brain regions support both the stability within and flexibility between task rules (Izquierdo, Brigman, Radke, Rudebeck, & Holmes, 2017; Uddin, 2021). However, despite this extensive research, translating findings to human models remains a challenge. This methodological gap hinders our ability to directly compare neural mechanisms underlying rule-switching across species. As a result, much remains unknown about how neurons, both within and across brain regions, coordinate to ensure both stability and flexibility in a rule-switching task. Critically, solving this gap requires a microstructural understanding of ensemble code transformations. Rather than a slow, gradual alignment, animal models demonstrate that the PFC mediates behavioural transitions during rule learning via abrupt, highly synchronized transitions between distinct neural ensemble states (Durstewitz, Vittoz, Floresco, & Seamans, 2010). Yet, capturing these rapid network state transitions remains exceptionally difficult in traditional setups.

Part of the reason for this is the absence of an ideal task to isolate the neural correlates for rule-guided decision-making process, and to gather a sufficiently large sample of such activities from low-firing prefrontal and anterior cingulate neurons. Moreover, there is currently no task in which rats can easily learn to switch multiple times between rules within a 30- to 60-minute session. Additionally, most rule-switch tasks do not contain a pre-response delay period in which the animal's posture remains neutral, which is critical for the identification of neural correlates for task rules. Without these parameters, researchers are often unable to distinguish between the neural activity required to maintain an internal rule and the activity associated with the physical preparation of a motor response. Moreover, the lack of high-frequency switches within a single session limits the statistical power needed to observe the subtle low-firing dynamics of the ACC and M2, leaving the precise mechanisms of cognitive flexibility obscured by behavioural noise and postural artifacts.

To optimally characterize the neural correlates for rule- and reward history-guided decision-making in rats and mice, a few conditions should be met: (1) A delay period preceding the motor response, when animals remain still and their postures are neutral, i.e. independent of the current rule and the upcoming choice (Duan, et al., 2021); (2) The trials are relatively short, so hundreds or more trials can be gathered in one session before satiation/fatigue sets in, during which activities of well-isolated single units are recorded continuously; (3) The rules are distinct but not too difficult, so the animals can learn to switch between them multiple times within a 30-60 minute session. If only one switch is achieved, in addition to the small number of correct responses and reduced sample of neuronal spikes, another issue is posed: rule-related changes in neural activities become confounded by time and other time-correlated factors, such as satiety, fatigue and boredom.

1.4 Our current task & goal

Based on the shortcomings of existing paradigms, we designed our current win-Stay/win-Shift Rule-Switching task (hereafter referred to as the "Rule-Switching task"). The Rule-Switching task differs from previous tasks as it provides no cue, only feedback, and is guided by working memory of the rule. With this, the Rule-Switching task allows the development of rule-switching expertise, complementing existing tasks in the behavioural neuroscience toolbox. Our rat model was specifically engineered to utilize a nose-fixated delay to isolate the maintenance of the internal rule from the motoric act of responding, effectively translating these clinical human concepts into a measurable circuit-level narrative in the rat. By forcing high-frequency switches within a single session over neutral delay intervals, this framework provides the statistical power necessary to observe the real-time population interactions of the ACC-M2 circuit as it coordinates macrostructural strategy bifurcations (Hamilton & Brigman, 2016) and captures discrete abrupt network state updates (Durstewitz, Vittoz, Floresco, & Seamans, 2010).

With our current task, we can explore the neural correlates underlying rule-switching behaviour. The main goal is to understand how frontal neurons encode rules and support rule-switching. To do so, we examined the behaviour of rats as they performed a novel decision-making task in which they must Shift or Stay. Shift means alternating sides (i.e., left and right) while Stay means repeatedly going to the same side. It is expected to find a training-induced enhancement in flexibility, characterized by a reduction in perseverance and reaction time, with an increase in accuracy over time. This study combines behavioural and single-unit analyses.

The subsequent chapters of this thesis systematically unpack the mechanics of rule-switching by moving from outward behavioural characterization to its underlying neural architecture. Chapter **2** establishes the behavioural baseline of this cognitive flexibility, quantifying the precise switch costs, error types, and motivational boundaries that emerge when subjects adapt to changing environmental rules. Having mapped these behavioural dynamics, Chapter **3** transitions to the underlying neural data, utilizing in-vivo recordings to capture the real-time ACC-M2 activity driving these choices. These two chapters funnel into a general discussion (Chapter **4**) that ties both worlds together. By pairing the behavioural profiles isolated in Chapter **2** with the precise physiological signatures uncovered in Chapter **3**, this thesis provides a high-resolution look at both the outward execution and the internal neural coordination of systemic rule-switching.

Chapter 2. Rats perform multiple switches on rule-switching task

2.1 Background

When navigating a Rule-Switching paradigm, a rat cannot rely on stable stimulus-response habits; instead, the animal must continuously compute two orthogonal variables: the spatial direction of the physical response (Side) and the active abstract behavioural framework (Rule). Investigating how the ACC and M2 concurrently track these variables offers a unique window into the neural mechanisms of flexibility, a process supported by evidence that specific cingulate-motor circuits dynamically update rule representations for sequential choice decisions (Takeuchi, et al., 2022).

A key challenge in decoding neural representations is understanding how behavioural strategy optimization modifies the underlying cognitive architecture over extended training. This influences both working memory capacity and the neural networks targeted by cognitive training (Constantinidis & Klingberg, 2016). Progression through the training stages—Beginner, Intermediate, Advanced, and Expert—of a rule-switching task involves a profound restructuring of behavioural outcomes, reflecting well-documented neural changes that occur after extensive training to perform cognitive tasks (Qi & Constantinidis, 2013). These stages were determined by the number of days that the animal had been training on the full task. Beginner is for days 1-2, Intermediate for days 6-7, Advanced for days 11-12, and Expert for days 16-17. With each day, the animal gained expertise, before mastery at 17 days. Early in training, the transition between rules is heavily disrupted by pro-active interference, where an animal blindly perseverates on a previous reinforced strategy (Birrell & Brown, 2000). As mastery is achieved, this rigid

perseveration drops significantly, yet animals begin to exhibit a distinct profile of regressive errors—briefly reverting to the old rule even after establishing a solid run of correct choices. This shifting error landscape implies that training does not simply eliminate mistakes, rather, it alters the internal stability and cross-interference of competing mental sets, driven by ongoing PFC plasticity during the learning of cognitive tasks (Tang, et al., 2022).

The primary objectives of this study were to evaluate how extended training shapes behavioural accuracy, reaction time, and error dynamics during the novel rule-switching task. Statistically, we aimed to identify the most parsimonious models that explain these changes. It was hypothesized that as rats gained experience and progressed from Beginner to Expert stages, (1) their overall choice accuracy would systematically increase, allowing them to respond correctly at a significantly higher frequency. Concurrently, it was expected that progressive training would lead to (2) a generalized reduction in reaction times, accelerating the animals' correct-response rates as a baseline cognitive adaptation independent of localized rule changes or spatial configurations. Regarding error profiles, it was hypothesized that (3) the count of perseverative errors would systematically decrease across training stages as the animals learn to adapt more rapidly to strategy shifts. Conversely, (4) regressive errors were anticipated to increase as animals achieve mastery, reflecting unique interference effects between highly familiar rules.

Finally, utilizing an information-theoretic framework for model selection (i.e., AIC), it is hypothesized that choice accuracy is best explained by a complex model incorporating individual fixed effects—trial Epoch (i.e., where a specific trial falls chronologically within the experiment), Rule type (Shift or Stay), training Stage (Beginner, Intermediate, Advanced, or Expert)—alongside a critical multiplicative interaction term between Rule type and Side (Left or

Right). This would demonstrate that accuracy is uniquely sensitive to the contextual interplay between cognitive rules and spatial orientation. Conversely, it was hypothesized that a highly restricted low-parameter model prioritizing training Stage as an isolated predictor would optimize the information criterion for reaction time, proving that physical processing speed accelerates as a generalized systemic adaptation over time.

2.2 Methods

2.2.1 Animals

Eleven, male Long-Evans rats (450-550 g) of 10-weeks were trained to perform the Rule-Switching task. Rats were housed in a colony with a 12-hour light/dark cycle, with all training and recording occurring during the light cycle.

All animal procedures were approved by the Central Commissie Dierproeven (CCD), The Netherlands, and conducted according to the Experiments on Animals Act (protocol codes, 2016-014-051).

2.2.2 Surgery

For detailed procedures please see our previous study (Ma, Hyman, Durstewitz, Phillips, & Seamans, 2016). Briefly, rats were anesthetized with 5% isoflurane, which was then gradually reduced to 2% for maintenance throughout the surgery. The head was placed in the stereotaxic frame and the scalp was shaved and disinfected. NSAID analgesic, antibiotic, and local

anesthetic were given before the incision. Following the incision, the fascia were removed and the skull was thoroughly cleaned, roughened, and several stainless-steel bone screws were installed. One craniotomy, approximately 3mm x 2mm, was created, followed by a durotomy. The guide tubes of the tetrodes were then placed on the exposed cortical surface. The implant housed two (2) bundles of tetrodes (Voigts, Newman, Wilson, & Harnett, 2020), one was positioned at AP = +3.2, ML = +0.5 from bregma to target the ACC, and the other at AP = +2.0, ML = +1.0 to target the M2. The durotomy was then protected with Vaseline before the entire implant was secured to the skull with dental cement.

2.2.3 Apparatus

Testing was conducted with rats placed in a specially designed polygon-shaped operant box (**Figure *1***), 33cm long and 30cm wide, constructed in-house with hard plastic boards and lined with white waterproof wallpaper on the interior. Animals have access to three 3cm-wide nose-poke ports: one square-shaped center port on the 10cm-wide front panel, and two circular side ports on the 15cm-wide left and right panels, respectively. Each port can be illuminated via a diffused blue LED on the back wall of the port and is equipped with an IR emitter and receiver to detect an entry of the animal's nose. A reward well was positioned in front of the center port where soy milk was delivered after each correct response, accompanied by an audible click of the solenoid.

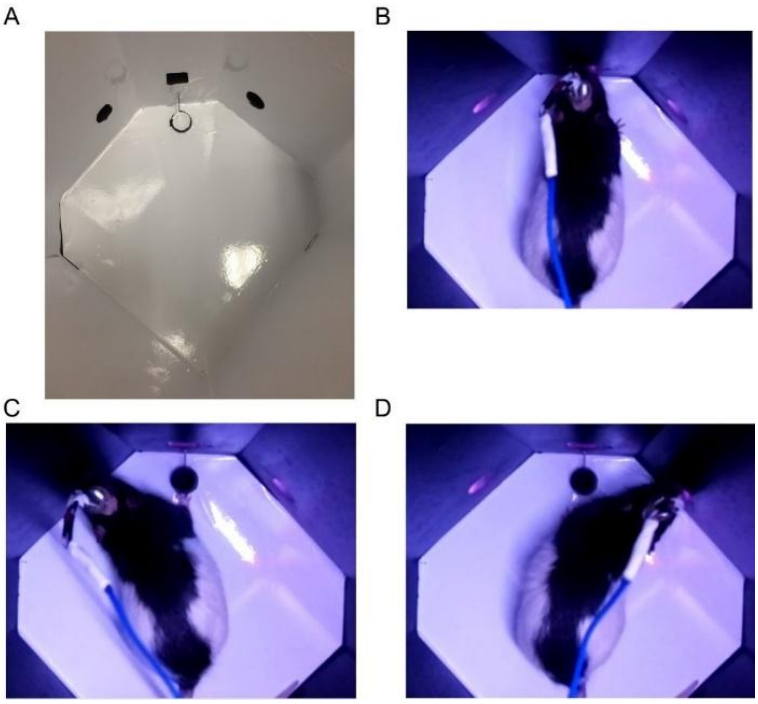


Fig. 1 | Operant box. Rats performed the Rule-Switching task from inside of a novel operant box. (**A**) The empty polygon-shaped operant box. Circular left and right nose-poke ports are seen on the slanted sides of the box. The square-shaped central port is seen on the horizontal edge. The circular white shape on the box floor, below the central port, is for soymilk reward output. The operant box with a rat inside equipped with tetrodes and doing a nose poke (**B**) in the center, (**C**) on the left-hand side, and (**D**) on the right-hand side.

2.2.4 Histology

Tissue preparation and sectioning

Following experimental endpoints, rats were deeply anesthetized using urethane intraperitoneal injection, and a 100-µA current was passed through the electrodes for 30 seconds.

Animals were then perfused with a solution containing 250 ml 10% buffered formalin (vol/vol), 10 ml glacial acetic acid, and 10 g of potassium ferrocyanide. This solution causes a Prussian blue reaction, which marks with blue the location of the iron particles deposited by passing current through the electrodes. The brains were then removed and stored in a 10% buffered formalin, 20% sucrose solution (wt/vol) for at least 1 week before being sliced and mounted to determine precise electrode locations. Given that multiple sessions were recorded from individual animals, the precise recording locations could not be derived from electrode lesions, but all electrode tracks were inferred between the entrance point and the dyed spot. Representative recording sites are shown in ***Supplementary figure 1***a,c and the ranges of recording are shown in ***Supplementary figure 1***b,d.

Microscopy and image analysis

Histological sections were imaged using a Leica DMI8 inverted microscope integrated into a Leica LAS X system (version 3.7.5.24914) configured with an AF 6000LX hardware platform. Transmitted light brightfield (TL-BF) tile scans were automatically acquired using a high-resolution Leica DMC4500 digital colour camera. For regional overview and tile-stitching, images were captured using a Leica HC PL FLUOTAR 10x dry objective lens with a numerical aperture (NA) of 0.30 and a constant field iris diaphragm setting of 21 and aperture diaphragm setting of 8.

Individual fields of view were acquired at a native resolution of 2560 x 1920 pixels without electronic binning, corresponding to an unbinned sensor bit depth of 12-bits per channel (RGB). Automated tile scanning was executed with an active 10% spatial overlap on both the X and Y axes to ensure optimal alignment during programmatic stitching. Complete panoramic mosaics of full histology sections were reconstructed via automated geometric alignment using

the integrated LAS X stitching algorithms. Final compiled stitched overview frames represented composite regions measuring up to 16,366 x 22,575 pixels, covering a physical area of approximately 7.95 millimeters x 10.97 millimeters. All acquisition parameters, including camera exposure (200.5 ms), light intensity (17%), and sensor gain (1x), were held constant across all sample extractions to permit normalized evaluation.

2.2.5 Task training

Pre-training steps

Rats progressed through a series of training steps to learn the final Rule-Switching task. Initially, a daily session contained 60 trials. As rats' experience grew, so did their motivation, and they were able to perform 120-180 trials within a daily session. All rats were first trained to nose-poke. During each trial, all three ports were illuminated simultaneously. Each nose-entry into any of the left, central, or right ports was rewarded, followed by an inter-trial interval (ITI) of two seconds, during which the LEDs were turned off. Trials were aborted after 10 seconds if no response was registered. Once the rats had obtained at least 10 rewards from each port, they progressed to the next training step, in which only one port was illuminated at a time. Since rats immediately developed a preference for one or two of the ports, this step ensured that the rats also learned to respond to less preferred ports. On the following sessions, the ports were illuminated in a pseudorandom order, one at a time. Rats must obtain at least 80% of the total reward before progressing to the next step. This training stage took up to eight sessions.

Trials started with the illumination of the central port. Initially, a nose-poke in the central port that lasted for 0.2 seconds extinguished its LED, triggered a white-noise buzzer, and was

immediately rewarded. Once the criterion was met, the delay was prolonged gradually at 0.1-second increments, until rats were able to 'nose-fixate' for one second in the central port. In delays longer than 0.4 seconds, small interruptions (i.e., less than 0.2 seconds) were allowed. Such interruptions represented a small adjustment in posture, without full retraction of their nose from the port. This training stage took up to eight sessions.

Following this step, in addition to completing the delay, upon exiting the central port, rats must also respond to the illuminated side port to obtain the reward. The order of side-port illumination alternated: if the target on Trial *t* was the left port, then the target would become the right port on Trial *t*+1. This stage took 3-10 sessions depending on the animal. Instead of illuminating only one side port as the target, both ports were illuminated upon the completion of the delay. Rats must respond in the same fashion from trial to trial as they did at the previous stage. This task was essentially a delayed alternation task, which required that rats hold the previous choice in their working memory, to guide their current choice. This training stage typically took only 3-4 sessions to complete.

Rule-Switching task

After animals received the implants and fully recovered from the surgery, they were re-trained on the Win-Shift Choice task. Once criteria were met, on the following days, they performed the Win-Shift/Win-Stay Rule-Switching task, while their neural activities were recorded. During each of the recording sessions, rats performed 200 or more trials. At each step, they must perform at 80% or above, excluding omitted trials but must not omit more than one third (1/3) of all trials, before proceeding to the next step. The recorded sessions were repeated

until the animals became highly proficient at the rule-switching task. In the recorded task, rats began with the Win-Shift rule. Once they performed correctly on eight (8) consecutive trials (excluding omitted trials), the rule would switch to Win-Stay. That is, they need to identify the correct side port and respond in this port in every trial. Thereafter, the rule would switch each time eight (8) correct trials were completed consecutively. See **Figure *2*** for a schematic representation and comparison of the rules.

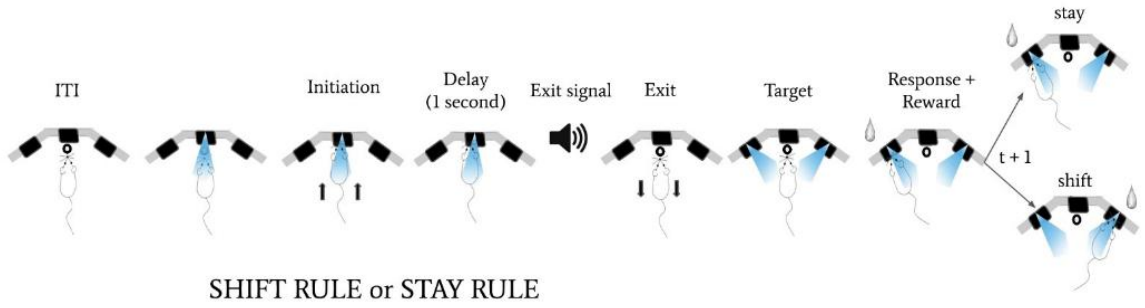


Fig. 2 | Schematic representation of the win-Shift versus the win-Stay task Rule. Rats performed the task, satisfying either the Shift or Stay rule based on which Rule they had previously completed. Three black boxes represent the ports (left, center, and right) through which rats made nose-pokes to log responses. Each port was equipped with a blue lightbulb. The circle beneath the center port represents the reward outlet (soymilk well) received for correct trials. Reward receipt is marked by liquid droplets. The arrows seen during the "initiation" and the "exit" phases represent the direction the rat is to move.

2.2.6 Data analyses

Two-factor ANOVA

To investigate the factors influencing performance, a two-factor ANOVA was employed for both accuracy and reaction time to evaluate the contributions of main effects and interactions between categorical experimental factors. Following this initial assessment, predictive modeling was conducted. Note that for reaction time in all analyses, only correct trials were used.

Logistic binary regression

For the accuracy analysis, a logistic binary regression model was utilized wherein the outcome variable *a* (accuracy) was dummy coded as 1 for a "correct" response and 0 for an "incorrect" response (**equation (*1*)**). We chose this model because Accuracy was a binary variable (0 = incorrect, 1 = correct), allowing us to model the probability of a correct response as a function of task-related variables. As the linear combination of the predictors—Stage (ST), Epoch (E), Side (S), and Rule (R)—increases, the probability of $P(a=1|x)$ approaches 1. Conversely, as the linear combination decreases, the probability approaches 0.

Multiple linear regression

Reaction time (RT) was modeled using a multiple linear regression approach (**equation *(2)***). This linear model included the intercept (β_0), an error term (ε) representing random noise, and the following predictors: Stage (ST), Epoch (E), Side (S), and Rule (R). $\beta_1, \beta_2, \beta_3$, and β_4 are the coefficients (weights) for the independent variables ST, E, S, and R, respectively. The

coefficients indicate how much the dependent variable of RT changes when the corresponding independent variable changes by one unit, assuming all other variables remain constant.

$$P(a = 1|x) = \frac{1}{1 + e^{-(\beta_0 + \beta_1 ST + \beta_2 E + \beta_3 S + \beta_4 R)}} \quad (1)$$

$$RT = \beta_0 + \beta_1 ST + \beta_2 E + \beta_3 S + \beta_4 R + \varepsilon \quad (2)$$

All-subsets regression

To identify the best-fitting models for both accuracy and reaction time, an all-subsets regression approach was implemented in Python using the *statsmodels* and *scikit-learn* libraries. This procedure consisted of three sequential steps: (1) feature selection, where potential predictors—Epoch, Rule, Stage, and Side—were isolated for analysis, (2) model iteration, where multiple regression models were systematically generated for every possible combination of these predictors, encompassing both main effects and two-way interactions, and (3) model comparison. The AIC was utilized as the primary metric for model selection (**equation *(3)***), where *k* is the number of estimated parameters and $\hat{L}$ is the maximum value of the likelihood function for the model. AIC rewards goodness-of-fit, while penalizing for the number of parameters to prevent overfitting; consequently, lower scores represent a better-fitting model (Zajic, 2022).

$$AIC = 2k - 2\ln(\hat{L}) \quad (3)$$

Higher-order factorial ANOVA

Finally, higher-order factorial ANOVAs were performed in *R* to conduct detailed error analyses, specifically examining the effects of Epoch, Stage, Rule, and Side on perseverative and regressive error counts. This higher-order factorial approach was selected because it permitted the simultaneous assessment of multiple categorical task variables and their interactions. For each of the two error types, the higher-order ANOVA was supplemented by a Tukey's honest significant difference (HSD) multiple comparisons of means post-hoc test, with significant differences recorded at a 95% family-wise confidence level.

2.3 Results

Dynamics of rule acquisition

All animals ($N = 7$ rats; IDs 7, 9, 10, 11, 12, 13, and 14) successfully mastered the rule-switching task within 6-12 weeks, yielding a total of 210 sessions (range: 6 to 58) containing at least one switch (i.e., at least one change from the Shift Rule to the Stay Rule or vice versa). Data from four rats (IDs 15, 16, 17, and 18) were omitted due to an insufficient number of completed sessions. Following the pre-training primarily on delay- and forced-choice Shift rules, animals achieved high performance within three days and learned to switch to the Stay rule within 1-2 days. Over 16 days of training, performance gradually increased and plateaued. Task performance was analyzed across four distinct training stages: Beginner (sessions D1 and D2), Intermediate (sessions D6 and D7), Advanced (sessions D11 and D12), and Expert (sessions D16 and D17).

The gradual improvement of performance accuracy

Rats rapidly learned to perform multiple intra-session rule-switches. Recovering to an accuracy above 80% took longer when switching from the Shift to the Stay rule than vice versa (

Figure .a versus

Figure .b). Overall, animals demonstrated higher accuracy during the Shift rule than the Stay rule (

Figure .c). A two-factor ANOVA revealed that learning Stage, post-switch Epoch, Rule, and Side (the latter as an interaction with Rule) all significantly contributed to performance accuracy. A subsequent logistic regression confirmed these findings ($F = 8.451$, $p < 0.05$, adjusted $R^2 = 0.011$), identifying significant main effects for Epoch, Rule, Side, and Stage, alongside a significant Rule x Side interaction (**Table *1***).

Table 1

Logistic regression testing for accuracy

Factor	***t***	***p*-value**	***β***
Epoch	4.678	2.99E-06*	0.074
Rule	-3.371	0.001*	-0.164
Side	-2.619	0.009*	-0.127
Stage	2.327	0.020*	0.016

Rule*Side	2.662	0.008*	0.082

Note. Asterisks denote statistical significance ($p < 0.05$).

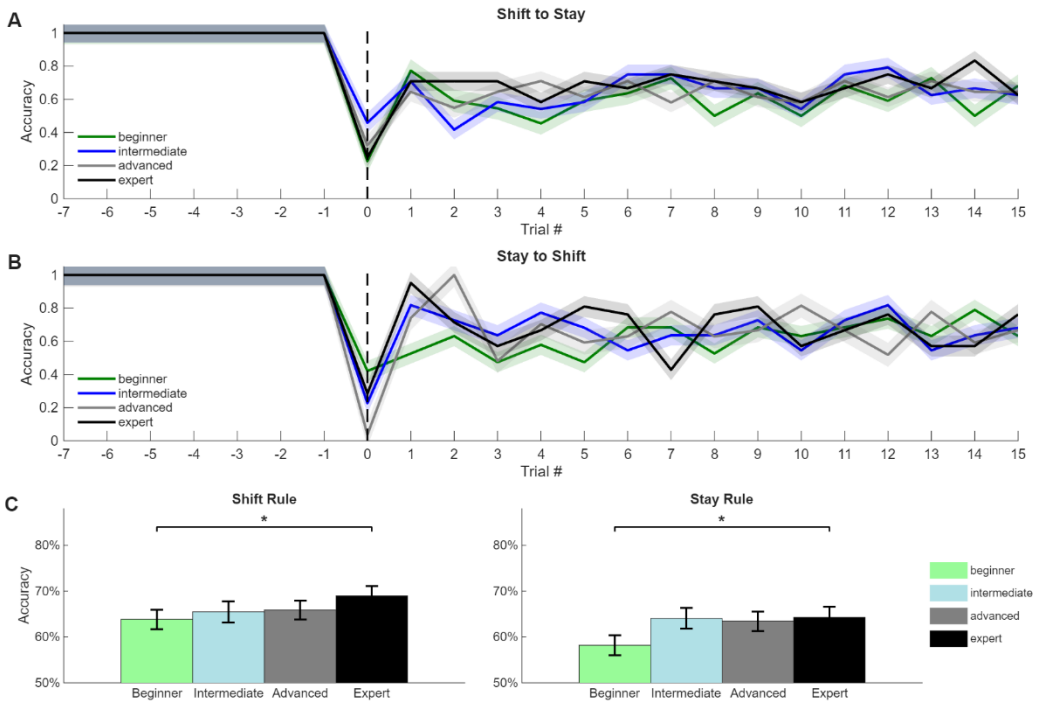


Fig. 3 | Accuracy averages. For both the Shift and Stay rule, accuracy significantly increased from Beginner to Expert. No other two stages saw significant differences. The Beginner, Intermediate, Advanced, and Expert stages are represented by green, blue, grey, and black, respectively. Choice accuracy after both (**A**) Shift-to-Stay switches and (**B**) Stay-to-Shift switches was significantly higher in the Expert stage than in the Beginner stage. The Advanced and Expert Stages yielded no significant differences, suggesting a plateau. (**A**) and (**B**) show sharp dips at x = 0 as 0 represents the first switch trial. Trial number on the x-axis range from 7 trials pre-switch to 15 trials post-switch. The shade around each line denotes standard error of the

mean (SEM). (**C**) Choice accuracy for both Stay and Shift during the four stages significantly increased from the Beginner stage to the Expert stage. However, with both rules combined, accuracy was insignificant for stage ($p = 0.07$).

Reaction time plateaus

Per the initial two-factor ANOVA, reaction time was predominantly modulated by experience on the rule-switching task (Stage). A multiple linear regression evaluated the continuous reaction time variable ($F = 219.3$, $p < 0.05$, adjusted $R^2 = 0.054$) and confirmed a highly significant main effect of Stage ($t = -14.810$, $p = 2.61E-48$; **Table *2***). No significant main effects or interactions were observed for Epoch, Rule, and Side.

Table 2

Multiple linear regression testing for reaction time

Factor	***t***	***p*-value**	***β***
Stage	-14.8	2.61E-48*	-0.105

Note. The asterisk denotes statistical significance ($p < 0.05$).

Overall, reaction times improved across trials as rats progressed from Beginner through Intermediate and Advanced stages to the Expert stage (**Figure *4***). A Tukey's post-hoc test conducted in *R* revealed that the reaction time relationship between the Advanced (Stage 3) and

Expert (Stage 4) phases was statistically non-significant ($p = 0.99$), indicating that the training effect on response latency had plateaued (**Figure *4***.c, d).

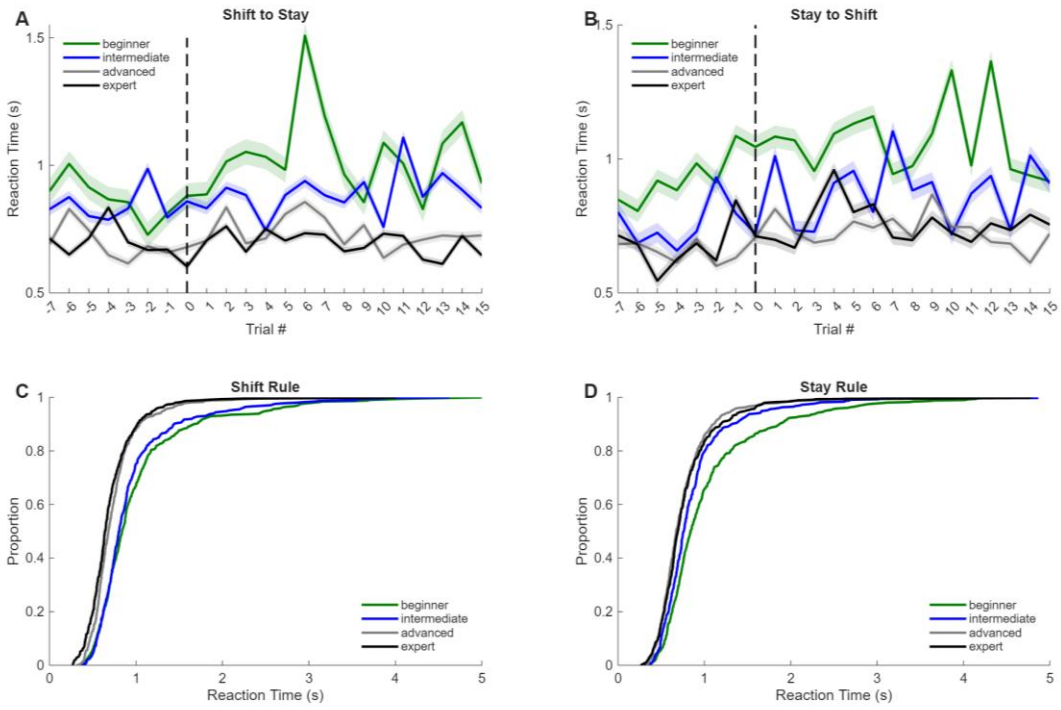


Fig. 4 | Reaction time averages. The four stages are Beginner (green), Intermediate (blue), Advanced (grey), and Expert (black). The top row shows reaction time amongst the stages when subjects performed (**A**) Shift-to-Stay and (**B**) Stay-to-Shift switches. (**A**) and (**B**) show that reaction time increases after a Rule switch, suggesting that the animal is taking more time to make their response. The dotted vertical lines represent where the Rule switch occurred (x = 0). The x-axes represent trial number, ranging from 7 trials pre-switch to 15 trials post-switch. Reaction time during both the (**C**) Shift Rule and (**D**) Stay Rule decreased as rats became more experienced. Advanced and Expert rats followed a strikingly similar trajectory, suggesting a

learning plateau. Each stage has a limit of x = 5 as, if rats had not logged a response within 5-seconds, the trial ended.

Rats demonstrate a shift from perseverative to regressive errors

Perseverative errors are when an error occurs immediately following an error (two 0s in a row). The higher-order factorial ANOVA error analysis for perseverative errors showed that both Rule ($F(1, 174) = 4.50$, $p = 0.04$) and Stage ($F(3, 174) = 4.02$, $p = 0.01$) had significant effects on the number of perseverative errors committed in a given session. The interaction between Rule and Stage did not yield a significant relationship ($p > 0.05$). Following the higher-order factorial ANOVA, we employed a Tukey's post-hoc test. The higher-order factorial ANOVA told us that there was a significant main effect of Stage and Rule, but not which levels differ, hence the post-hoc test. These factors have multiple levels and therefore necessitate pairwise comparisons. According to mean differences (*MD*) from the Tukey's post-hoc test, the Stay rule had more perseverative errors than the Shift rule ($MD = 0.44$, $p = 0.04$; ***Figure 5***). For Stage, differences were only significant between the Beginner and later stages. The Beginner stage had significantly more perseverative errors than the Intermediate stage ($MD = 0.87$, $p = 0.03$), the Advanced stage ($MD = 0.77$, $p = 0.04$), and the Expert stage ($MD = 0.96$, $p = 0.01$; ***Figure 5***). There was no additional reduction in perseverative error from intermediate to later stages.

Regressive errors are when an error occurs after having three consecutive correct responses (three 1s, followed by a 0). For regressive errors, we used the same higher-order factorial ANOVA but this time modeling the effects of Rule, Stage, Epoch, and Side on regressive error count. The error analysis for regressive errors showed that only Stage had a significant effect on the number of regressive errors committed in each session, $F(3, 174) = 3.03$, $p = 0.03$. Additionally, we performed a Tukey's post-hoc test for regressive errors because Stage

has multiple levels and we needed to ascertain between what levels the significance occurred. There was only a significant *MD* in regressive error count between the Beginner and Expert stages (*MD* = 0.47, *p* = 0.03; ***Figure 5***). The number of regressive errors increased from the Beginner to the Expert stage. Rule and the interaction between rule and stage did not yield a significant relationship (*p* > 0.05).

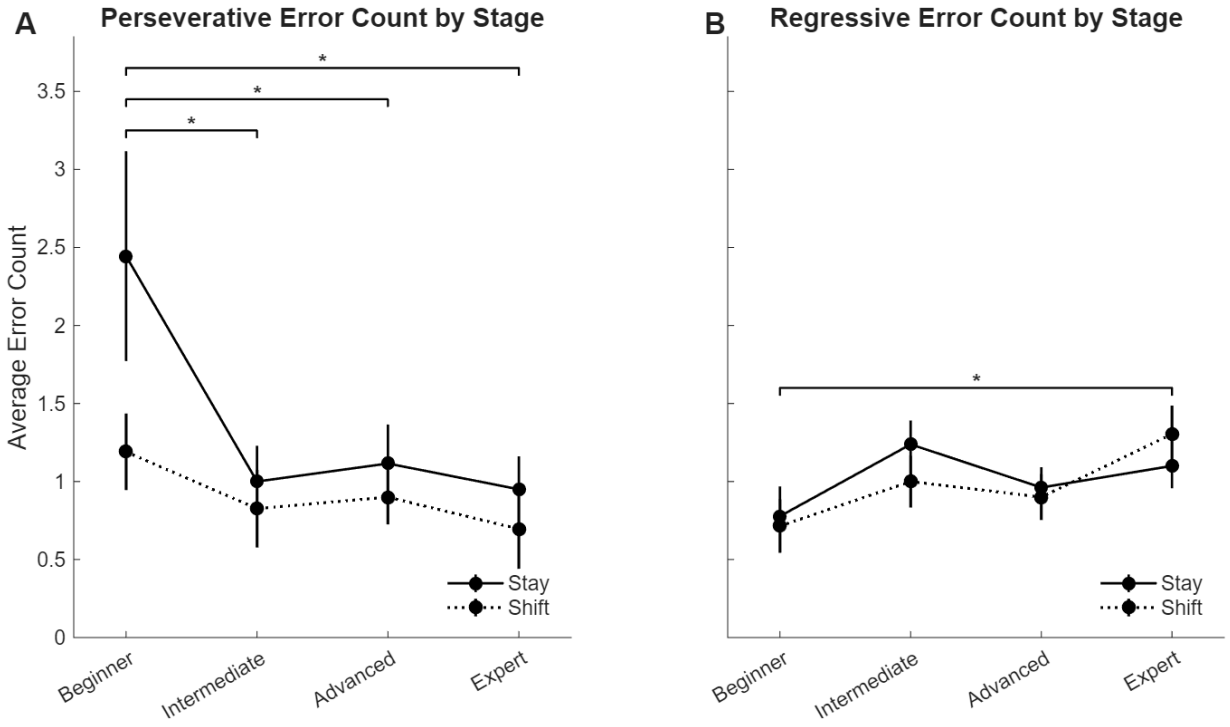


Fig. 5 | Perseverative and regressive error count by stage. The Stay Rule is represented by the solid line while the Shift Rule is represented by the dotted line. Average error count is calculated from the first 15 trials post-switch of sessions 1 and 2 (Beginner), 6 and 7 (Intermediate), 11 and 12 (Advanced), and 16 and 17 (Expert). (**A**) Overall perseverative error across Rule decreased from Beginner to later stages. (**B**) Overall regressive error count across Rule increased across Stages.

AIC minimization reveals distinct predictors for performance accuracy and reaction time

All different regression combinations were calculated to identify the equation that yielded the lowest AIC possible. AIC scores determine which regression model is best for a given data set, with lower scores meaning the model fits the data better (Zajic, 2022). This was calculated using Python. For accuracy, the lowest AIC was attained when all factors, except for Side, were individually implicated. Side was then included as an interaction term to yield even lower AIC (**Figure *6***). For reaction time, the lowest AIC was attained when only Stage was implicated (**Figure *7***).

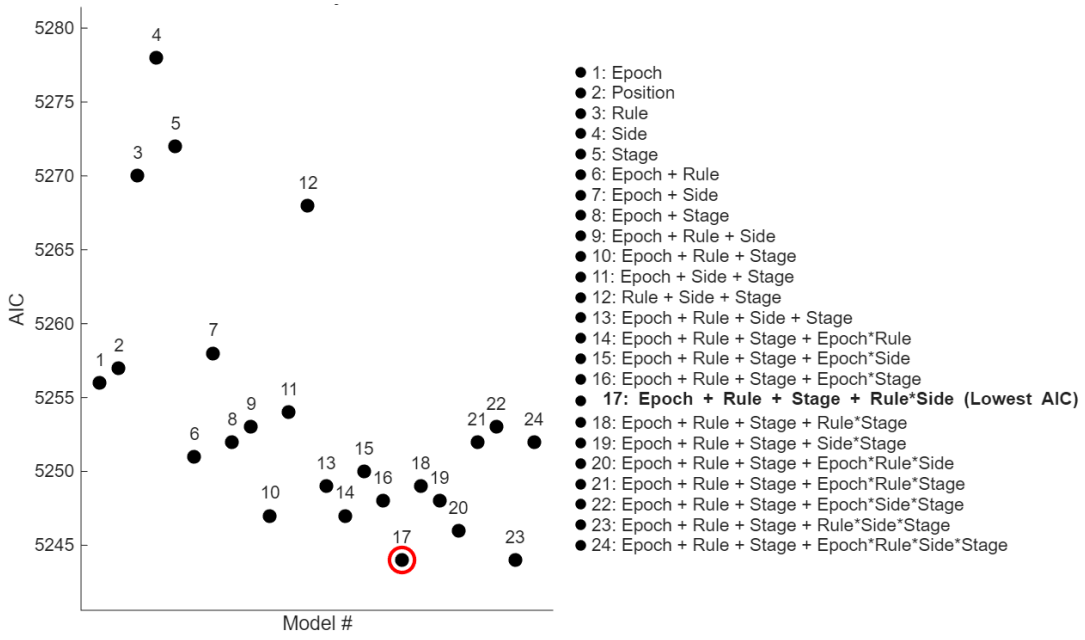


Fig. 6 | Accuracy AIC results. Lowest AIC was obtained at model #17 with Epoch + Rule + Stage + Rule*Side.

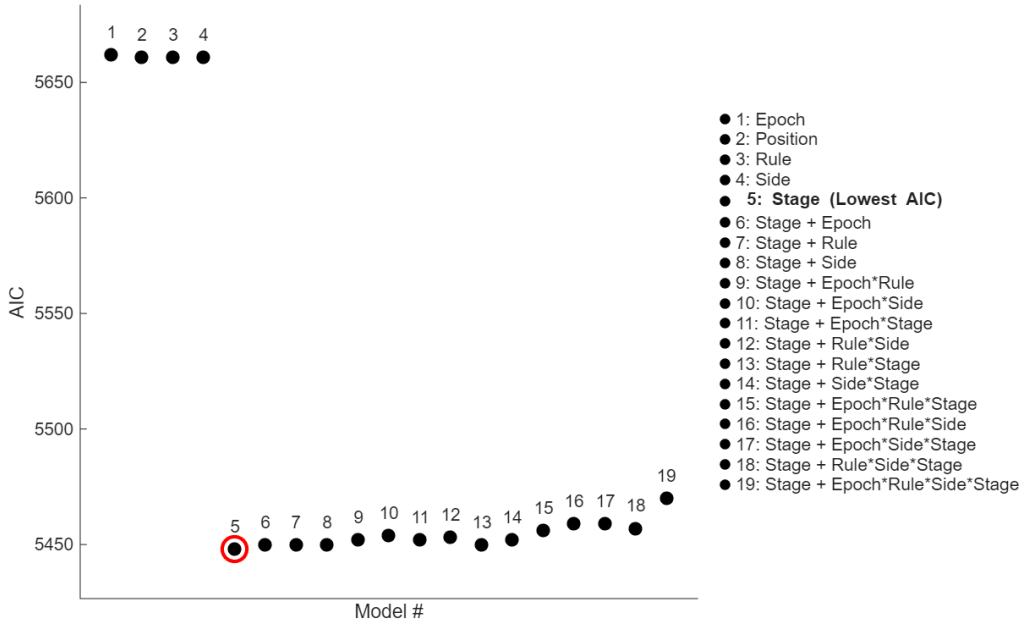


Fig. 7 | Reaction time AIC results. Lowest AIC was obtained at model #5 with Stage.

2.4 Discussion

The findings of this study demonstrate that training on the rule-switching task induces multi-staged optimization of both accuracy and reaction time. As animals progressed from the Beginner to the Expert stage, their ability to execute strategic shifts matured significantly, accompanied by a structured reorganization of their error profiles. These results provide insights into the cognitive architectures governing proactive control, interference management, and behavioural automation.

2.4.1 The evolution of the rat error landscape

The systematic reduction in perseverative errors from the Beginner stage onward indicates that experience primarily enhances an animal's proactive cognitive control. Early in training, strategy transitions are heavily disrupted by proactive interference, wherein the previously reinforced rule remains dominant and suppresses flexibility (Braver, 2012). The dramatic drop in consecutive errors following a rule-switch demonstrates that training allows rats to actively suppress outdated strategies and rapidly update their internal behavioural goals in response to changed environmental contingencies.

Intriguingly, task mastery did not eliminate mistakes entirely; rather, it introduced a distinct profile of regressive errors during the Expert stage. In cognitive psychology, regressive errors—where an individual briefly reverts to prior strategy after a baseline string of successful actions—are an indicator of competition between highly familiar internal mental sets (Monsell, 2003). As both rules become deeply ingrained in the animal's repertoire, they exist in a state of mutual activation. A regressive error represents a transient failure to maintain the active rule's dominance over a competing alternative strategy, rather than a generalized attention deficit. A regressive error is a specific failure of inhibition, while a generalized attention deficit is a failure of engagement (Birrell & Brown, 2000; Floresco, Block, & Tse, 2008; Milner, 1963).

In executive functioning testing (e.g., WCST, Stroop Task) researchers differentiate regressive errors due to competing strategy from attention deficits by looking at behavioural patterns, error timing, and underlying neural mechanisms. Regressive errors (competing strategy) see the subject successfully acquiring and applying a new rule for several trials but suddenly reverting to a previously learned rule. This error is non-random, following the exact logic of the alternative rule. Attention deficit errors are erratic and non-systematic errors where the subject

makes off-target mistakes, misses cues, responds prematurely, or fails to stick to any strategy consistently. Reaction time spiking is another useful differentiator for the meaning of these mistakes. In regressive errors, reaction time often spikes right before or during the error as the brain attempts to resolve interference between the Shift and Stay rules. Once prompted or corrected, the individual usually snaps back to the active rule quickly. In a general attention deficit, reaction time fluctuates wildly throughout, and performance degrades gradually over time rather than in sharp, rule-bound lapses. Thus, in the case of attentional deficits, feedback does not reliably stabilize performance. Reaction time spiking will be analyzed in future analysis.

2.4.2 Information-theoretic dissociation of performance metrics

The application of the AIC for model selection revealed a powerful mathematical dissociation between the factors driving choice accuracy and those driving reaction time. Choice accuracy was best explained by a highly parameterized model that crucially included a multiplicative interaction between Rule and Side. This indicates that a rat's accuracy is not a simple function of abstract rule execution; instead, it is highly sensitive to how concrete spatial variables, such as Side, intersect with cognitive Rules. This spatial-cognitive interaction suggests that the brain maps rules directly onto directional motor frameworks to guide successful choices. Conversely, the information criterion for reaction time was optimized by a restricted minimal model isolating training Stage as the sole meaningful predictor. The total exclusion of Rule, Side, and trial Epoch from the optimal reaction time model demonstrates that the acceleration of reaction time during correct trials across weeks of training is a generalized, systematic adaptation.

Chapter 3. Learning-induced population recruitment in ACC-M2 networks

3.1 Background

In complex decision-making tasks, the ACC and M2 are critical for integrating past outcomes, or feedback, with preparing future motor actions, or preparation (Alexander & Brown, 2019). In fact, the coordination between the ACC and M2 is a cornerstone of flexible behaviour. Current literature suggests that the ACC is specialized for evaluating task effort and processing outcomes (Kolling, Behrens, Wittmann, & Rushworth, 2016), while M2 is primarily associated with action selection and the preparation of motor sequences (Erlich, Bialek, & Brody, 2011). While it is well-established that these regions encode task-relevant variables, it remains unclear how these representations evolve as an animal transitions from a Beginner ("Early") to an expert ("Late") stage. A central debate in systems neuroscience is whether learning sharpens existing neural tuning or recruits a larger ensemble of neurons to ensure robust task execution (Driscoll, Pettit, Minderer, Chettih, & Harvey, 2017).

During rule-guided decision-making, population activity within the frontal cortex often displays high degrees of mixed selectivity—where individual single units simultaneously encode multiple, overlapping task parameters (Rigotti, et al., 2013). While mixed selectivity expands the computational capacity of a network, it also complicates traditional analytical interpretations at both the single-neuron and macro-population levels.

Standard dimensionality reduction techniques, such as Principal Component Analysis (PCA), collapse global variance into dominant trajectories but lack the mathematical framework to isolate the distinct contributions of independent experimental variables. Demixed Principal

Component Analysis (dPCA) resolves this constraint by partitioning population variance into mathematically orthogonal marginalizations, allowing for the isolated tracking of specific task features across time while maintaining the benefits of dimensionality reduction (Kobak, et al., 2016).

In the context of the rat ACC-M2 network, the ACC and M2 form a critical functional circuit. Navigating a Rule-Switching task forces this circuit to concurrently process and sustain two primary, orthogonal task variables: the spatial response profile (Side) and the active cognitive constraint (Rule). Applying dPCA to neural population data recorded during distinct behavioural windows—specifically the Delay period of internal maintenance and the active Response period of motor execution—provides a structured mechanism to explore how these parameters are distributed across the ACC-M2 network. Crucially, these low-dimensional population state-space trajectories can be utilized in a targeted, top-down approach to rank individual cells by their component loadings, isolating highly leveraged exemplar units for validation using raster plots and peri-stimulus time histograms (PSTHs).

A prevailing theory in systems neuroscience is the *Efficient Coding hypothesis*, which suggests that learning should sharpen neural tuning, reducing the metabolic cost by using fewer but more precise neurons to represent a variable (Barlow, 1961). Conversely, the dynamic neural recruitment framework suggests that expertise engages ensembles of neurons that pass the baton of information over time, ensuring that task rules are resistant to noise and maintained over long delays (Friedman & Robbins, 2022; Miller & Cohen, 2001). To adjudicate between these competing frameworks, a high-resolution, folded Area Under the Receiver Operating Characteristic (auROC) decoding analysis can be implemented to measure longitudinal shifts in informational intensity versus informational density across training stages (Early—first two days

of training; Late—last two days of training). Furthermore, the ITI—the period containing Feedback from a previous trial and Preparation for the next—is a critical window for this adaptation. In rats, the ACC-M2 circuit has been shown to support rule-based switching, but how the temporal persistence of these rules changes with expertise is poorly understood. The current study utilizes high-resolution auROC decoding to measure these subtle shifts in population dynamics across weeks of training.

The primary objectives of this chapter are to (1) characterize the low-dimensional population dynamics of the ACC-M2 circuit during rule-switching, isolating the single-unit profiles driving these states, and (2) determine whether learning-driven expertise is achieved via single-neuron sharpening or population-level expansion. Essentially asking how groups of neurons in the ACC and M2 region work together to track different parts of a task (the physical Side chosen versus the abstract Rule followed)?

Regarding learning-driven circuit optimization, it was hypothesized that the ACC-M2 network adapts to behavioural expertise through a population-level expansion of task representation rather than single-unit refinement (Rigotti, et al., 2013). Specifically, the informational intensity (Mean Strength) of individual task-responsive units is expected to remain stable or unchanged across progressive training stages, whereas the informational density (the percentage of significant decoding bins, or PercSig) would expand significantly in Late compared to Early training Stage.

Finally, we hypothesized that recruiting this neural population creates a lasting memory trace (temporal integration) across the ACC-M2 circuit. This means that: (1) The rich information encoded during the task does not fade during the quiet inter-trial delays; it stays strongly and consistently active. (2) This sustained activity connects two critical phases: looking

back to evaluate the outcome of the last trial and looking forward to plan the movement for the next one.

3.2 Methods

For apparatus, animal, surgery, and task-related procedures see 2.2.

3.2.1 Spike sorting

Automated template-matching spike sorting was performed on the broadband, bandpass-filtered extracellular data using *Kilosort 4* (Pachitariu, Sridhar, Pennington, & Stringer, 2024). This algorithm tracking approach allowed for the mathematical isolation of overlapping spike waveforms and corrected for intra-session electrode drift by continuously tracking non-stationary spatial distributions of templates across the probe matrix. Kilosort results were then manually verified by trained experimenters.

3.2.2 Data analyses

Demixed principal component analyses (dPCAs)

Neural recordings were obtained from 209 neurons in ACC and 203 neurons in M2 during performance of the rule-switching task. Population activity was analyzed separately for the Delay and Response periods, resulting in four independent demixed principal component analyses (dPCAs) as follows: (1) ACC Delay, (2) ACC Response, (3) M2 Delay, and (4) M2 Response. For each analysis, neural activity was organized according to the two task variables:

Side and Rule. dPCA was then applied to separate neural population variance into four marginalizations: Side, Rule, Side x Rule interaction, Condition-independent activity. Condition-independent activity captured temporal patterns shared across all task conditions, whereas the remaining marginalizations isolated variance associated with specific experimental factors.

The resulting demixed components provided a low-dimensional representation of neural population activity across time. Population trajectories and component loadings were examined to identify dominant patterns associated with each task variable. To facilitate subsequent single-neuron analyses, neurons were ranked according to their loading on selected dPCA components. To visualize the activity pattern in the group, neurons within the top 10% of absolute loading values were selected as dominant contributors. This percentile threshold is consistent with established feature-selection frameworks in high-dimensional neural data (Cunningham & Yu, 2014; Velten, et al., 2022). These neurons were then selected as exemplar neurons for visualization using raster plots and PSTHs.

Folded Area Under the Receiver Operated Characteristic (auROC) analysis

To evaluate how learning-related expertise alters neural representations across the frontal cortex, a comprehensive database of 12,264 unique neuron-comparisons was compiled into a long-format array. Neural activity was recorded from the ACC and M2 of eleven rats during the win-Shift/win-Stay, rule-based decision task. We utilized a folded Area Under the Receiver Operating Characteristic (auROC) metric to quantify how well individual neurons discriminated between task variables (Side and Rule). Our first metric, Mean Strength, was calculated as the average folding auROC across specified time bins to measure signal intensity. The second metric, Percent Significant (PercSig), is defined as the percentage of 100 millisecond bins within a window that reached statistical significance ($\alpha = 0.05$) via permutation testing. Next, temporal

epoch analyses targeted the Feedback window (2-3 seconds before the switch) and the Preparation window (1-2 seconds before the switch) relative to the trial onset.

Linear mixed-effects (LME) model

Given that multiple neurons were sampled from individual subjects across sequential training sessions, treating individual single units as independent observations would violate basic assumptions of variance independence and artificially inflate statistical power. To control for this nested, hierarchical data structure, we employed Linear Mixed-Effects (LME) Models fit by Maximum Likelihood (ML). The models tested the fixed effects of Stage (Early vs. Late), Region (ACC vs. M2), Epoch (Feedback vs. Preparation), and Comparison Type (Side vs. Rule). In all formulated models, subject identity (RatID; 11 levels) was designed as a random effect intercept, effectively partitioning inter-animal variance from the fixed experimental effects. The final LME looked at 5/11 RatIDs, all of which possessed selective neurons based on mean signal strength (MeanStrength)—the average of significantly selective neurons. For signal strength calculation, we included neurons with at least 1 time bin in which significant coding for a task parameter was achieved. The five rats were the result of a filter which omitted sessions with less than three rule-switches. A minimum of three switches guaranteed the comparison between at least one condition (i.e., Rule or Side).

Two-way interaction model

To test whether ACC and M2 optimize task representations via individual unit precision (sharpening) or population-level recruitment, we first fit baseline two-way interaction models using training Stage (i.e., Early or Late) and cortical Region (i.e., ACC or M2) as predictors. The

initial fixed-effects structure evaluated the interaction between training Stage and cortical Region on individual neuron discriminative capacity using the formula seen in **equation (*4*)**. Later, to assess population informational density (PercSig), we performed modeling via the fixed-effects framework seen in **equation (*5*)**.

$$MeanStrength \sim 1 + Stage \times Region + (1 \mid RatID) \tag{4}$$

$$PercSig \sim 1 + Stage \times Region + (1 \mid RatID) \tag{5}$$

3.3 Results

Exploratory dPCA reveals overlapping task representations in ACC and M2 with minimal qualitative dissociation

The application of dPCA to Delay-period activity revealed representations associated with Side, Rule, and condition-independent activity in both ACC and M2. The temporal structure of the population trajectories indicated that neural activity evolved systematically across the Delay period and that multiple task variables contributed to the observed population dynamics. This analysis was used to identify individual neuron activity, determining where in the Delay-period the activity was ramping up. The dPCA was especially useful for this Rule-Switching paradigm because of the use of internal representations. When it comes to making the correct

choice, educated decisions are based on information from the previous trial. We cannot possibly know when signals would rise without the application of the dPCA which shows patterns in time.

In ACC Delay, condition-independent activity accounted for 54% of the explained variance, followed by Side-related activity (27%), the Side x Rule interaction (11%), and Rule-related activity (8%; Top-Left of ***Figure 8***). In a dPCA, total percent variance explained quantifies the proportion of the dataset's total variability that is captured by the selected dynamic principal components (PCs). This serves as a metric for how well the reduced-dimensional model preserves the overall informational content and time-dependent structures of the original time series. The dominant condition-independent components exhibited gradual temporal changes across the Delay period that was largely shared across all task conditions. These components may reflect population-level processes related to task timing, anticipation, or general cognitive engagement rather than specific task variables. Meanwhile, Side-selective representations were evident. In ACC, the strongest Side-related component (PC2) explained 19.1% of the total variance and clearly separated left- and right-Side trials.

In M2 Delay, condition-independent activity accounted for 61% of the explained variance, followed by Side-related activity (17%), the Side x Rule interaction (12%), and Rule-related activity (10%; Bottom-Left of ***Figure 8***). M2 Delay showed similar patterns to ACC Delay, although the corresponding Side component (PC2) explained a smaller proportion of variance (10.4%). Additional Side-related components accounted for smaller amounts of variance in both regions.

Overall, the delay period analyses revealed remarkably similar population-level organization in ACC and M2. In both regions, neural activity was dominated by condition-

independent temporal dynamics, with Side information representing the largest task-related source of variance and Rule-related activity contributing a comparatively smaller proportion.

In ACC Response, condition-independent activity accounted for 45% of the explained variance, followed by Side-related activity (48%), the Side x Rule interaction (4%), and Rule-related activity (3%; Top-Right of ***Figure 8***). The strongest Side-selective components showed clear divergence between left- and right-side responses, consistent with increased encoding of behavioural choice during this epoch. Meanwhile, in M2 Response, condition-independent activity accounted for 47% of the explained variance, followed by Side-related activity (39%), the Side x Rule interaction (7%), and Rule-related activity (7%; Bottom-Right of ***Figure 8***). Similar to the ACC region, the strongest Side-selective components showed clear divergence between left- and right-side responses, consistent with increased encoding of movement execution.

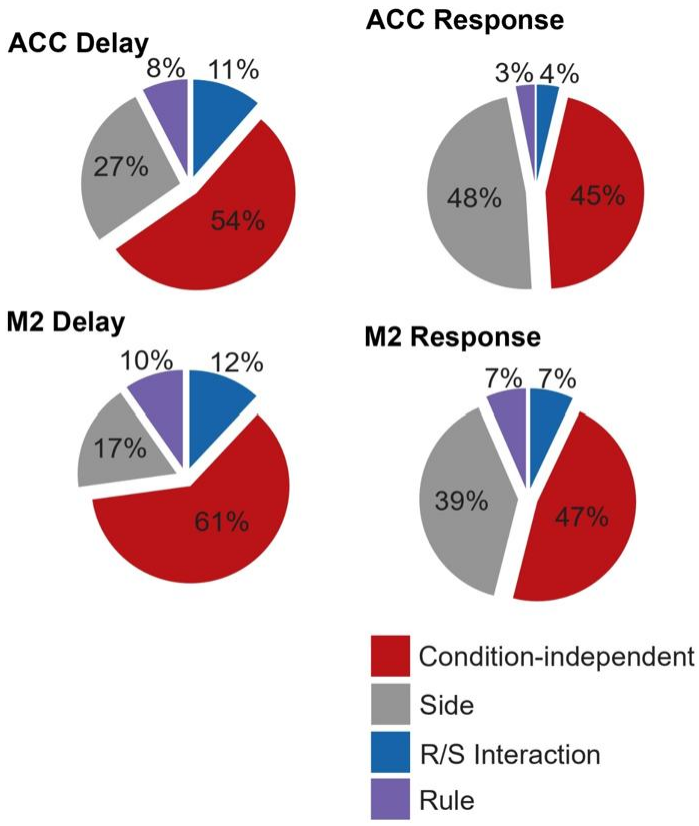


Fig. 8 | Demixed principal component analysis (dPCA) explained variance by condition. We captured two temporal windows: Delay (1-second before the Delay end) and Response (1-second before the Response onset). The Left Column corresponds to the Delay-period, while the Right Column corresponds to the Response-period. We recorded from both the ACC and M2 regions. The Top Row corresponds to the ACC, while the Bottom Row corresponds to M2. The red colour represents condition-independent explained variance. The conditions of Side, Rule-Side Interaction, and Rule are seen in grey, blue, and purple, respectively. A higher percentage means that condition explains more variance.

The peri-Response period revealed a similar organizational structure to the Delay period. Condition-independent activity remained the largest source of explained variance in both

regions. However, Side-related activity increased substantially during the Response period relative to the Delay epoch. The strongest Side-selective Response components were highly similar between ACC and M2. Both regions displayed pronounced separation of left- and right-side trials as animals approached response execution. Condition-independent Response components showed nearly identical temporal profiles in the two regions, characterized by strong changes in population activity preceding response onset. These shared temporal dynamics further suggest that both regions participated in similar population-level processes during behavioural execution.

Qualitatively, the overall organization of population activity appeared similar between ACC and M2. Both regions exhibited components associated with Side and Rule information, and neither region demonstrated a clearly distinct pattern of representation that would suggest fundamentally different coding schemes during the Delay epoch. Instead, the results suggested that both cortical areas participated in representing task-relevant information during this period. The interaction marginalization accounted for a smaller portion of the observable structure relative to the primary task variables, although interaction-related activity was present in both regions. A similar pattern emerged during the Response period. Neural populations in both ACC and M2 displayed separable components associated with Side, Rule, and condition-independent activity. Population trajectories reflected changes in neural activity as animals executed behavioural responses, while still maintaining representations related to task conditions.

As observed during the Delay epoch, qualitative comparisons between ACC and M2 revealed substantial similarities. The dominant features of the population activity were broadly conserved across regions, suggesting overlapping representation of task variables during response execution. Although minor differences were present in individual components and

temporal trajectories, no major qualitative dissociation between ACC and M2 was apparent from the exploratory dPCA analyses.

The dPCA analyses did however provide a principled method for identifying neurons that contributed strongly to specific population-level representations. Neurons were ranked according to their component loadings, and those within the top 10% of loading values (~21 neurons in each region) were selected for further investigation. Examination of raster plots and PSTHs from these high-loading neurons revealed examples of neural activity patterns that closely reflected the population-level structure identified by dPCA. These exemplar neurons served as representative examples of how task-related information was encoded within individual cells while contributing to broader population dynamics (***Figure 9***).

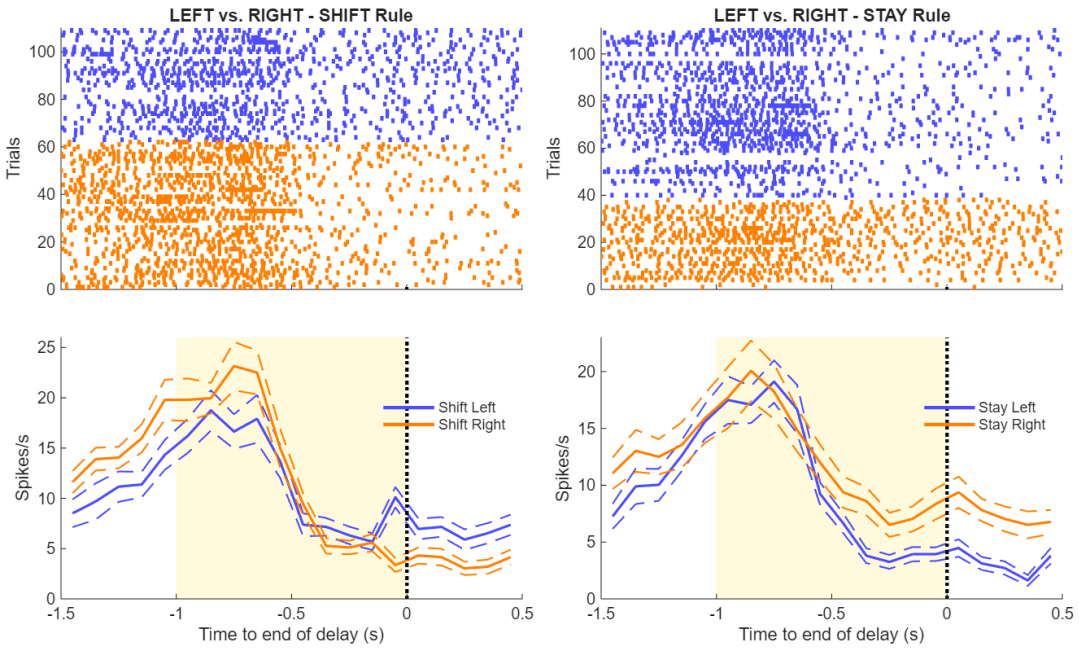


Fig. 9 | Exemplar neuron ACC Delay. Exemplar neuron from animal ID14, from date 2022-01-12 and time 17-46-48 for Neuron #2. All four panels represent the same animal, day, and neuron.

The top row are raster plots, with the left representing the Shift Rule's Left side versus Right side and the right representing the Stay Rule's Left side versus Right side. The bottom row PSTHs correspond to the raster plot directly above them. Overall, the four panels represent a similar pattern which shows increased firing rates right before and during the highlighted delay period (i.e., from -1.5 to ~-0.75 seconds), before a substantial decrease in activity ~0.5 seconds before the end of the delay period.

Learning drives parallel informational density expansions in ACC and M2 populations rather than a singular localized hub

Our analytical narrative was anchored to two primary dependent metrics computed across 100 millisecond bins. Firstly, Mean Strength, a measure of informational intensity quantified via the average folded auROC across specified time windows. This model analyzed a sample size of $N = 7{,}336$. Secondly, we had PercSig, a measure of informational density and temporal persistence. Percent Significant is quantified as the percentage of bins within an experimental window that achieved statistical significance ($\alpha = 0.05$) via permutation testing ($N = 12{,}264$ observations). We ran the auROC on all 11 rats, however, of these 11, 5 met criteria of having at least 10 trials per completed contrast (contrast refers to which Rule and Side combinations were being compared; there are a total of four different contrasts) for both an Early and Late stage. None of the omitted animals had a Late Stage, as they did not complete enough trials.

Evaluation of the percent selective neurons revealed a highly significant, global main effect of learning stage ($\beta_{\text{Stage_Late}} = 8.1925$, $SE = 0.676$, $t_{(12260)} = 12.117$, $p = 1.34 \times 10^{-33}$). This represents the total proportional density of statistically significant encoding windows across the 12,264-observation framework. During Early acquisition (Sessions 1-5), the baseline population representation was compressed (Intercept = 12.37% ± 1.84%). Upon reaching baseline mastery

for Late acquisition (Sessions 10-14), the network expanded its active encoding envelope by an absolute value of +8.19%, extending across both ACC and M2. A standard two-way interaction model (Stage * Region) confirmed that this informational expansion was universally shared between structural nodes; the baseline density of the M2 did not deviate significantly from the ACC ($\beta_{Region_M2} = 1.0505$, $p = 0.178$), and its tracking trajectory across learning remained bound to the global trend ($\beta_{Stage_Late:Region_M2} = -0.0108$, $p = 0.991$).

As seen in ***Figure 10*** (top panels), in the Early training phase, the baseline proportion of task-selective neurons in the ACC stood at an estimated 12.37%. By the Late training phase this population footprint expanded significantly, experiencing an absolute surge of +8.19% to reach a final density of 20.56% ($p = 1.34 \times 10^{-33}$). This recruitment pattern was similar within the M2, which tracked the ACC's surge with an influx of selective units, establishing a population density similar to the ACC ($p = 0.991$ for the stage-by-region interaction term). This uniform expansion demonstrates that learning does not compress or favor a singular localized hub. Instead, it systematically broadens the network manifold, driving a greater population recruitment cooperatively across both ACC-M2 regions.

Global Stage x Epoch interactions mask asymmetric longitudinal adaptation specific to motor cortex populations

Following these analyses and with $N = 11$ rats, we evaluated the spatial capacity of network encoding by assessing the percentage of active, threshold-passing time intervals per unit (PercSig). The primary two-way model demonstrated a main effect of Stage ($\beta = 2.630$, $SE = 0.531$, $t(12260) = 4.951$, $p = 7.48 \times 10^{-7}$). Across both structures, the proportion of task-informative time intervals expanded universally, yielding an average increase of approximately 2.63% to 3.18% more active encoding bins per neuron in Late sessions.

While the omnibus, global interaction did not clear the global significance threshold (Stage by Region interaction, $p = 0.146$), direct pairwise tracking reveals that the mathematical weight of the longitudinal shift is driven by cellular recruitment in M2. When dissecting population densities via the three-way framework (Stage*Region*CompType; CompType being the four different comparison types of Rule versus Side), there was no significant main effect of CompType ($F(1, 12256) = 3.39$, $p = 0.0655$), This indicates that the temporal bandwidth occupied by Rule and Side did not reliably differ at the population level.

Individual single-unit tuning intensity remains functionally static across training

In contrast to the dramatic expansion seen at the population level, analyzing the Decoding Strength (*auROC* magnitude tracking across 7,336 threshold-passing, significant data bins) demonstrated that single-unit coding intensity remains flat. The Early baseline coding intercept settled tightly at $auROC = 0.5582 \pm 0.002$. The fixed-effect coefficient for Late-Stage mastery (Expert) was significant ($\beta_{Stage_Late} = -0.00436$, $SE = 0.0014$, $t_{(7332)} = -3.16$, $p = 0.0016$). This drop was marginally mitigated in the M2 by a minor interaction term ($\beta_{Stage_Late:Region_M2} = 0.0047$, $SE = 0.0019$), which brought the Late M2 encoding trace back to a net-neutral baseline ($0.5582 – 0.0042 + 0.0047 \approx 0.5587$).

As seen in ***Figure 10*** (bottom panels), functional preservation and invariance of single-unit tuning intensity was stark in contrast to the massive inflation of the active population pool. The informational intensity of individual localized processing units—quantified via single-neuron decoding strength (MeanStrength auROC)—showed statistical shifts that were very small in effect size. The mean decoding strength for individual task-modulated ACC neurons hovered close to a baseline auROC of 0.5582 during Early sessions. Upon reaching task expertise (i.e., Late sessions), individual tuning curves registered a small decline of -0.0043 ($p = 0.0016$).

Individual M2 units exhibited a complementary, marginal upward shift of +0.0046 (p = 0.015), functionally nullifying the baseline change observed in the ACC. Kernel density estimations of the data distribution clouds highlight that the underlying probability profiles of individual single-unit auROCs are virtually superimposable between Early and Late stages. The structural center of mass for both regional networks remains pinned near baseline thresholds.

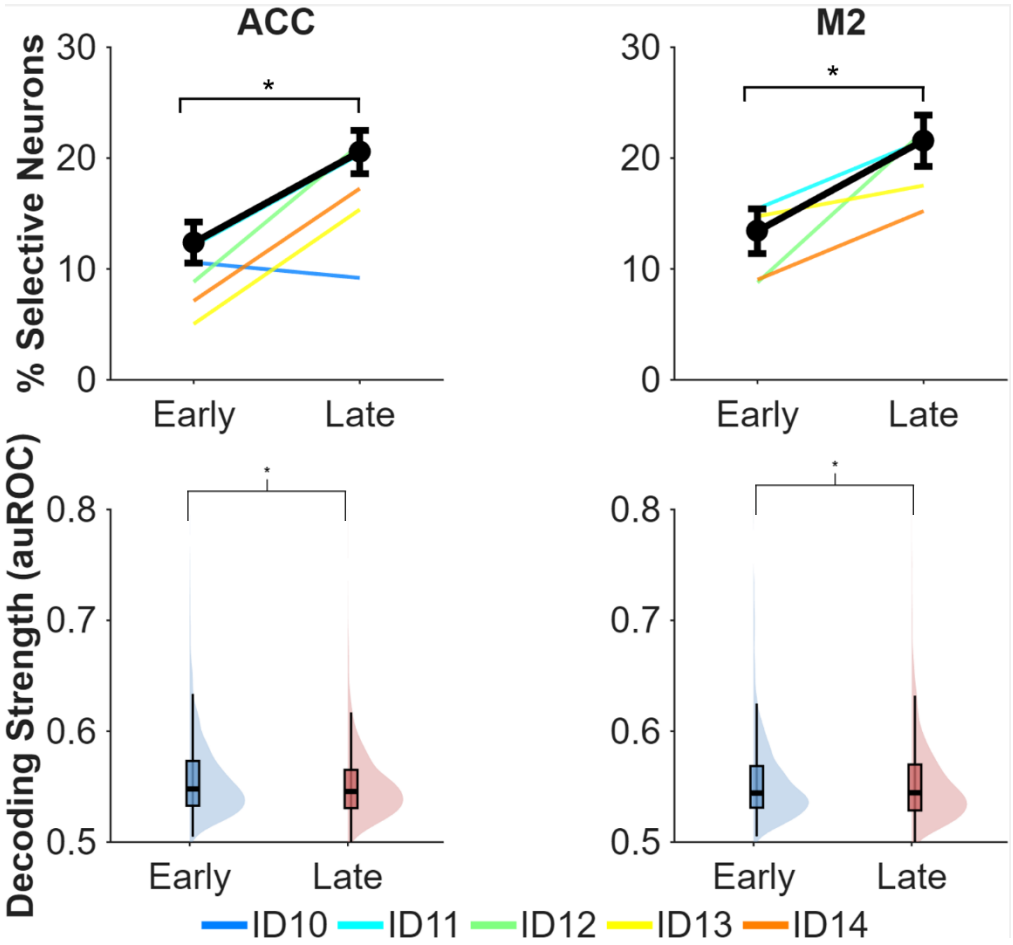


Fig. 10 | Percent selective neurons and decoding strength (auROC) for ACC and M2 regions in Early versus Late sessions. Learning-induced expansion of ACC-M2 population

manifolds preserves single-unit tuning dynamics. Top Row: Longitudinal tracking of population informational density (measured as the percentage of task-selective neurons per session; % Selective Neurons) across Early (Sessions 1-5) and Late (Sessions 10-14) training blocks for the ACC (left panel) and M2 (right panel). Coloured lines track empirical session means for individual subjects (n = 5 rats; ID10-ID14). Overlaying bold black markers and connecting vectors represent global population fixed-effects estimates (±SEM) derived from the LME regressions. Both regions exhibit a highly significant, parallel recruitment surge (~ 8.2% absolute expansion, $p < 0.001$). Bottom Row: Structural distribution of informational intensity (quantified as single-neuron decoding strength via auROC) for ACC (left panel) and M2 (right panel). Shaded asymmetric violin plots represent kernel density probability estimations mapping the global variance footprint of single-neuron observations across Early (blue clouds) and Late (red clouds) stages. Superimposed isolated markers depict true population fixed-effects means (±SEM), demonstrating a functional preservation of absolute cellular tuning intensity despite massive network recruitment. All source calculations control for nested subject-specific intercept variances (1|RatID).

Asymmetric expansion in the ACC-M2 circuit is condition-dependent and globally invariant to inter-trial epoch gating

To determine if this population expansion was uniform across different dimensions of the behavioural task, we evaluated a 3-way interaction model (Stage * Region * CompType) for the percentage of Selective Neurons. Marginal Type-III ANOVA testing identified significant individual main effects for learning Stage ($F_{(1, 12256)} = 99.24$, $p = 2.74 \times 10^{-23}$) and comparison architecture ($F_{(1, 12256)} = 5.94$, $p = 0.015$). Crucially, the analysis revealed a significant 3-way interaction between training Stage, recorded Region, and comparison type (Rule or Side

comparison; $F_{(1, 12256)} = 3.91$, $p = 0.048$; $\beta_{\text{CompType:Region_M2:CompType_Side}} = 3.56\% \pm 1.80\%$, $p = 0.048$). This interaction indicates that the geometry of the ACC-M2 manifold does not expand symmetrically; instead, the network selectively modulates its spatial density based on the specific cognitive Rule being processed within each anatomical Region.

Finally, an LME auditing temporal processing windows (Feedback vs. Prep epochs) demonstrated that while population recruitment is highly dynamic across the ITI ($F_{(1, 12256)} = 6.94$, $p = 0.0084$) for the main effect of Epoch, this temporal gating operates independently of the learning process. No significant interaction terms were found for Stage:Epoch ($p = 0.232$) or Stage:Region:Epoch ($p = 0.285$).

Learning amplifies execution-level tuning magnitude in M2 but not ACC

The analysis with $n = 5$ rats containing all trial types was done as exploration. These exploratory results gave hint of potential effect, thus, to ascertain the results with $n = 5$, the code was expanded and the omitted animals (without Late stages) were added to a second, repeat analysis. This analysis contained $N = 11$ animals, increasing the number of animals for the Early stage only. At $N = 11$ the analysis included more sessions and yielded greater statistical power. To evaluate whether the absolute magnitude of single-unit encoding changes as a function of task experience when $N = 11$, we constructed the same LME model; predicting MeanStrength using fixed effects of learning Stage (Early versus Late), anatomical region (ACC versus M2), and their interaction, with random intercepts nested by subject (RatID). The analysis revealed a significant two-way interaction between Stage and Region ($\beta = 0.00379$, $SE = 0.00189$, $t(7140) = 2.008$, $p = 0.0447$).

Post-hoc examinations of the fixed effects indicated that within the ACC, individual single-unit tuning magnitude remained remarkably stable across longitudinal development (β_{Stage_Late} = 0.00017, p = 0.901). Conversely, single-units within M2 demonstrated an enhancement in their raw representation strength over time. This region-specific plasticity is further reflected in the standardized stage effect sizes: while the ACC tracking profile underwent a negligible change (0.0055 Standard Deviations), M2 individual single-unit tuning magnitude expanded by 0.1083 Standard Deviations. To test whether this tuning amplification was preferential for Rule or Side parameters, we executed a three-way model, including CompType (i.e., Rule versus Side). The three-way interaction (Stage:Region:CompType) failed to clear the significance threshold (p = 0.455), establishing that the longitudinal tuning amplification observed in M2 represents a generalized enhancement across both Rules and Side.

Longitudinal scaling modulates inter-trial epoch gating to dynamically reorganize pre- and post-choice representations

To understand the temporal distribution of these representations during the ITI, we evaluated coding densities across distinct task epochs: post-choice evaluation (Feedback; -3 to -2 seconds) and action preparation (Preparation; -2 to -1 seconds). A marginal ANOVA test revealed a prominent main effect of Epoch (F(1, 12256) = 15.01, p = 0.0001), confirming distinct baseline densities between the two phases. Most notably, the analysis yielded a significant Stage by Epoch interaction (F(1, 12256) = 4.95, p = 0.0262; **Figure *11***). This interaction demonstrates that the expansion of population coding density is not uniformly distributed over time; instead, as animals transition from Beginner performance to proficiency, the temporal profile of active representations undergoes an organization shift that alters the relative density of informative bins during outcome evaluation versus motor preparation.

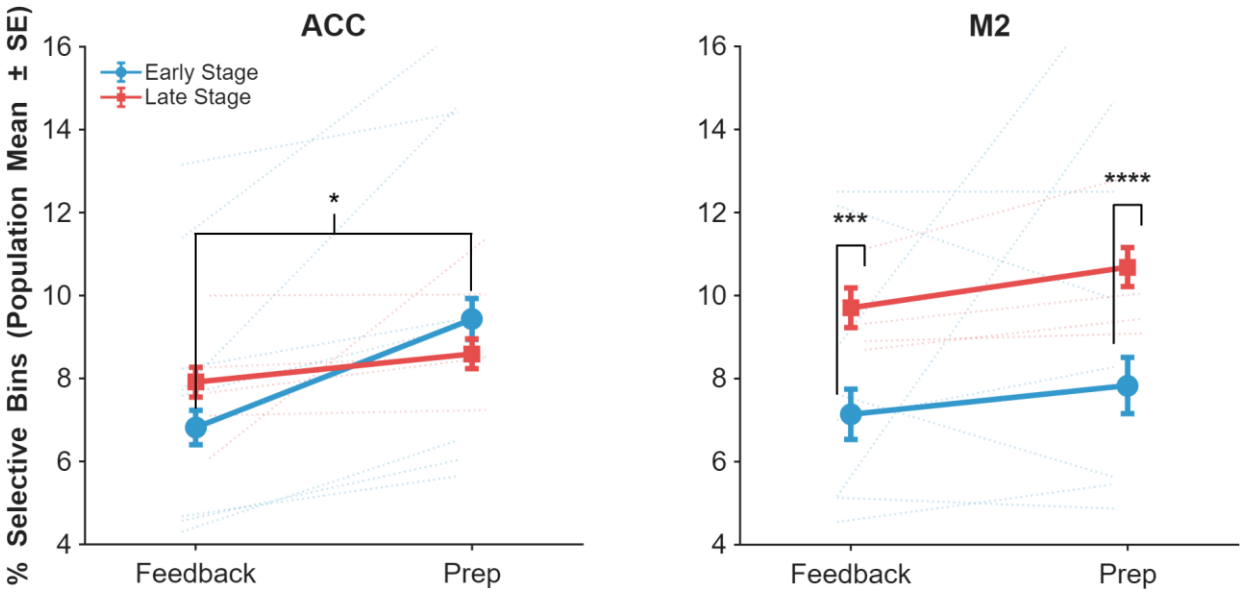


Fig. 11 | Reorganization of task information density across learning stages and inter-trial epochs. Network information density (expressed as the mean percentage of task-selective time bins per neuron; PercSig) is plotted across two distinct epochs of the ITI: post choice outcome evaluation (Feedback, -3 to -2 seconds relative to the trial onset) and action preparation (Preparation, -2 to -1 seconds). Solid circles, squares, and heavy lines indicate the global population mean ±SEM calculated from the LME model. Light dotted lines denote the underlying longitudinal trajectories of individual animal subjects across the group ($N = 11$ rats). The triple asterisk for M2 Feedback window indicates significance of $p < 0.001$, while the quadruple asterisk for M2 Preparation window for $p < 0.0001$. There is a significant global Stage x Epoch interaction ($p = 0.0262$), revealing a computational reorganization of network timing over learning. In the Beginner state (Early Stage; blue lines), encoding capacity expands sharply as the trial approaches execution. Following task mastery in the Expert state (Late Stage; red lines), this temporal gating is systematically restricted: baseline tracking capacity shifts upward

during the feedback monitoring window, leading to a compressed, stable processing volume leading into the subsequent trial. The flattening of the trajectory is a global feature where this parallel temporal adaptation occurs independently of anatomical division, running globally across both the ACC (left panel) and M2 (right panel). However, the underlying baseline shift is anchored heavily by M2's expansion.

Post-hoc analysis of interaction terms

Unpacking our global interaction model using post-hoc pairwise contrasts (Holm-Bonferroni corrected) reveals that the experience-dependent optimization of this ACC-M2 circuit is heavily asymmetric, driven predominantly by a massive longitudinal expansion within the M2. The ACC maintains a statistically stable volume of active information density across longitudinal training (Early versus Late, $p > 0.09$). Expert animals recruit a drastically higher density of selective processing windows within M2 during both the Feedback window (adjusted $p = 0.00011$) and the Preparation window (adjusted $p = 0.000022$).

3.4 Discussion

The neural recordings spanning the ACC and M2 reveal how ACC-M2 population manifolds and single-unit tuning restructure to support flexible decision-making. By deploying dPCA and high-resolution auROC decoding, this chapter highlighted a highly integrated spatial-cognitive circuit and provided definitive evidence regarding the structural network adaptations that accompany behavioural expertise.

3.4.1 Overlapping latent spaces and sensorimotor gating

Applying dPCA untangled mixed selectivity characteristics of frontal neurons (Rigotti, et al., 2013), isolating independent variance streams for Side and Rule. The striking qualitative similarities in population state-space trajectories between the ACC and M2 suggest that these regions communicate rapidly and extensively during rule-switching, rather than as strictly segregated modules. This cross-regional alignment fits perfectly within the multiple-demand system framework, which posits that complex executive control relies on distributed, highly interconnected cortical networks rather than isolated regional specialization (Duncan, 2010).

3.4.2 Population expansion and dynamic neural recruitment over metabolic extremes

If learning depended solely on a uniform, network-wide sharpening of the accuracy of individual single units, we would see a clear, positive upward trajectory in all our MeanStrength curves. This would reflect enhanced signal-to-noise ratios or intensified metabolic allocation within dedicated Expert neurons. Instead, individual single-unit decoding performance within the ACC remained functionally stable across longitudinal development. Interestingly, a clear division of labour emerged at this single-cell level: while the ACC acts as a steady, upstream corporate manager of task constraints, M2 exhibits targeted, plasticity, significantly amplifying the raw tuning magnitude of its units over time. This selective single-unit sharpening in M2 points to its dynamic role for well-learned behavioural routines (Barthas & Kwan, 2017).

At the population level, the circuit optimizes its broader processing pipelines by expanding its active temporal capacity, recruiting a significant global increase in selective encoding windows across the ITI. The mathematical weight of this longitudinal shift is driven

predominantly by cellular recruitment within M2, while the ACC population remains statistically stable. This dramatic, epoch-specific expansion within the motor cortex suggests that the ACC-M2 circuit prioritizes regional population recruitment over single-cell intensification to support Expert-state execution. By distributing Rule and Side information across a wider, more integrated ensemble matrix, the ACC-M2 circuit gains two clear advantages being (1) robustness to noise and (2) metabolic subsidization. For robustness to noise, a distributed population architecture is far less vulnerable to the stochastic firing fluctuations or loss of individual single neurons (Zylberberg, Pouget, Latham, & Shea-Brown, 2017). For metabolic subsidization, spreading the computational load across a larger pool of cells prevents individual units from needing to maintain exhausting, high-frequency metabolic firing rates (Chugani, 1999; Huttenlocher, 1984).

For the sub-population of rats ($n = 5$), the presence of a significant 3-way interaction (Stage:Region:CompType, $p = 0.0479$) reveals that this manifold expansion is not a generic, unguided increase in baseline activity. Instead, it represents a structured reconfiguration of the network's geometry. The positive interaction coefficient ($\beta = 3.55\%$) demonstrates that Late-Stage recruitment of M2 is heavily customized to support egocentric action coordinates (CompType_Side). This aligns well with the known anatomical role of M2 in mapping motor planning pathways (Barthas & Kwan, 2017). Meanwhile, the ACC maintains balanced coverage across abstract Rules and Side. This confirms that as cognitive expertise develops, the ACC-M2 circuit selectively channels specific task variables into their appropriate functional nodes.

Ultimately, these results contribute to our understanding of how cognitive control skill development shapes the ACC-M2 circuit. The transition from a Beginner to an Expert is not characterized by a uniform, network-wide intensification of single neurons firing harder. Instead,

while execution nodes like M2 selectively sharpen their individual output, the broader network optimizes by learning to coordinate and open more parallel processing channels across the collective population manifold. This population expansion creates a stable and distributed neural workspace capable of flexible, rule-guided behaviour.

Chapter 4. General discussion

In the present study, we used a novel operant box to observe cognitive flexibility in rats during a rule-switching win-Shift/win-Stay task. To our knowledge, and as of the publication date of this thesis, this represents the first successful implementation of multiple within-session rule switches in an operant box that incorporates a neutral-posture delay period. The overall results of the study support the primary hypotheses, demonstrating that reaction times and the number of general errors significantly decreased while accuracy increased as training and level of experience progressed. This reflects both a training-induced enhancement in cognitive flexibility and an optimized trial-and-error strategy during rule transitions. Specifically, the rats demonstrated the ability to execute multiple Shift and Stay switches within a single 60-minute session. This behavioural efficiency yielded a high-density trial structure that is essential for gathering the high-volume neural data required to analyze low-firing ACC and M2 populations.

Our neural analyses confirmed our hypotheses regarding experience-dependent circuit optimization, revealing a clear distinction between localized cellular features and global network scaling across the full $N = 11$ rat population. By deploying dPCA across the group, we successfully untangled mixed selectivity into independent latent spaces, revealing a highly integrated, shared spatial-cognitive blueprint across the ACC and M2. LME modeling of this 11-animal framework unlocked the underlying cellular, geometric, and temporal dynamics driving long-term network adaptation. At the single-unit level, individual neural drive (MeanStrength) revealed a clear division of labour: while individual ACC units remained remarkably stable and unchanged over weeks ($p = 0.901$) to serve as a steady cognitive scaffold, downstream M2 units

exhibited targeted, execution-level plasticity by significantly amplifying their raw tuning magnitudes over time ($p = 0.0447$).

At the population level, the circuit appears to support rule-switching skill development via an asymmetrical process. While collapsing across individual task variables revealed a highly significant global main effect of Stage across the full framework ($p = 7.48 \times 10^{-7}$), post-hoc pairwise contrasts unmasked a highly localized transformation. M2 drove this longitudinal expansion by drastically increasing its selective encoding windows (PercSig), while the ACC population envelope remained statistically stable ($p > 0.09$). Within this population expansion, the network implemented a structured geometric and temporal reconfiguration. In the full group, a significant Stage x Epoch interaction ($p = 0.0262$) demonstrated a profound rebalancing of network timing. While Beginner animals relied on a reactive strategy—task-related information was weak prior to a sharp rise right before the next trial—expert networks completely flattened this trajectory, front-loading task information directly into the post-choice feedback window to maintain a proactive, always-ready processing state.

Our paradigm's relevance in RDoC construct and subconstructs

The present study sits within the National Institute of Mental Health's (NIMH) Research Domain Criteria (RDoC) framework. Within the RDoC, the current work concerns the *Cognitive Systems* domain. Specifically, we care about *Cognitive Flexibility*, hereafter referred to in this thesis as "flexibility" for short. Flexibility is the ability to adapt behaviour in response to changing rules (Brady & Floresco, 2015). As per extensive research, flexibility is commonly impaired across multiple psychiatric conditions, such as schizophrenia, ADHD, and mood

disorders (Cochereau, Zandvoort, Santariius, & Mandonnet, 2021; Milton, et al., 2021; Murphy, et al., 1999; Murphy, Michael, & Sahakian, 2012; Perpiñá, Segura, & Sánchez-Reales, 2017; Rady, Elsheshai, el Wafa, & Elkholy, 2012; Wang, Hu, & Li, 2022; Zastrow, et al., 2009).

The rule-switching task used in this study operationalizes cognitive flexibility in a controlled experimental setting. This allowed for multiple lines of analysis emphasized by the RDoC framework. Firstly, we have the behavioural level with performance accuracy, switching success, and learning rate. Second, the neural systems level with population-level encoding of task variables (i.e., Side, Rule, and their interaction). By examining how both behaviour and neural representations evolve with training, this study aligns with the RDoC goal of linking observable behaviour to underlying neural mechanisms across time. Importantly, the identification of learning-dependent changes in neural encoding, particularly in interaction terms, such as Side x Rule, provides insight into how cognitive control processes are dynamically constructed rather than statically represented. Early-stage variability in neural activity may reflect the formation of task-relevant representations. Meanwhile, later-stage stabilization may indicate more efficient or automated processing. Furthermore, the use of dimensional learning stages (Beginner, Intermediate, Advanced, Expert) parallels the RDoC emphasis on continuous variation rather than categorical distinctions. Rather than classifying behaviour as simply impaired or intact, this approach captures gradual improvements in cognitive flexibility and their neural correlates.

Overall, this work contributes to the RDoC framework by (1) providing a model of learning-dependent changes in cognitive control, (2) demonstrating how population-level neural dynamics relate to behavioural adaptation, and (3) supporting a multi-level, dimensional approach to understanding executive functioning. Consequently, these findings may inform translational

models of psychiatric dysfunction, where deficits in cognitive flexibility can be understood as disruptions in the formation, stability, or efficiency of neural representations underlying rule-guided behaviour.

4.1 Statistical predictors of task success: Trial Epoch, Rule type, and Stage

When examining accuracy across the training timeline the GLM evaluated individual and interactive contributions, identifying Epoch, Rule, and Stage as significant predictors of task success/accuracy. Rather than operating as isolated variables, these three components function as a dynamic predictive matrix. Each factor exposes a distinct adaptation within the ACC-M2 circuit as the animal moves from Beginner to Expert automation.

Rule type as a behavioural filter

As established by our GLM model, Rule type serves as a primary baseline predictor of performance accuracy across the entire longitudinal timeline. A steady, significant drop in accuracy whenever a block transitioned from win-Shift to win-Stay demonstrates that these rules do not carry equivalent cognitive weights. From an evolutionary perspective, as aforementioned, win-Shift matches the animal's default, unmanipulated spatial foraging script. Win-Stay, by contrast, forces a structural friction into the decision-making apparatus: the animal must consciously overwrite its innate behavioural momentum. Statistically, this makes Rule type a powerful filter for isolating raw executive effort. The persistent performance gap between the two rules, even in highly trained Expert animals, proves that evolutionary cognitive biases are not truly erased. Rather, these biases are continuously managed by top-down ACC-M2 architectures that require ongoing computational bandwidth to suppress instinctual impulses.

Trial Epoch as a temporal axis of adaptation

Unlike in later models discussed in this section, in this behavioural framework trial Epoch does not denote isolated ITI windows, but rather trial position. Thus, trial Epoch is the chronological placement of a trial within a given block or session (e.g., immediate post-switch trials). Operating as a continuous temporal axis, trial Epoch serves as a critical predictor of how efficiently the circuit clears proactive behavioural interference and updates its rules.

In early learning stages (Beginner), trial Epoch yields a highly volatile predictive profile. Immediately following a rule switch (early Epoch positions), accuracy drops steeply as the animal commits perserverative errors, continuing to execute the historical rule. As the trial Epoch advances deeper into the block, accuracy gradually recovers, mapping the slow trial-and-error nature of the animal's reactive strategy. As expertise develops (Expert), the predictive weight of trial Epoch undergoes a transformation. Expert animals minimize the performance drop at early block positions, efficiency matches our population-level temporal flattening: by converting the ITI into an always-active, proactive processing state, the Expert network can clear interference and update its internal models within a single trial, flattening the predictive impact of trial Epoch on performance.

Learning Stage as a dimensional trajectory

Rather than treating learning as a binary switch (e.g., pre- versus post-reversal), our framework uses learning Stage (Beginner, Intermediate, Advanced, Expert) as a continuous, multi-tiered predictor. This dimensional approach successfully captured non-linear configurations within the circuit that categorical models inherently mask.

In the initial stages, the predictive value of Stage tracks the rapid suppression of perseverative errors. This reflects the primary clearing of behavioural interference as the ACC-

M2 network learns to detect a rule change across early trial Epochs. In mid-to-late stages, however, the predictive trajectory of Stage shifts its focus from error reduction to structural optimization. While accuracy begins to hit a behavioural plateau and reaction times floors out due to musculoskeletal constraints, Stage continues to predict the expansion of the neural population manifold.

This late-stage predictive trajectory provides empirical evidence for the Adaptive Coding Hypothesis. It demonstrates that expertise is not a static behavioural endpoint where learning ceases (Duncan, 2001). Rather, it is characterized by an ongoing internal rebalancing of the network. Even when behavioural readouts stabilize across rules and trial Epochs, advancing learning Stages continue to predict an asymmetric influx of selective encoding windows (Shibata, et al., 2017). This fundamentally reconfigures the ACC-M2 workspace into dynamic neural recruitment; a metabolically subsidized, and noise-resistant engine designed to protect Expert-state execution.

4.2 Reaching the cognitive ceiling: Reaction time and practice

In contrast to the multi-variable complexity driving accuracy, our analysis of response latencies revealed that the training Stage emerged as the sole significant predictor of reaction time. As animals transitioned from naïve Beginners to proficient Experts, their reaction times underwent a steep, non-linear decline. This optimization eventually reached plateau between the Advanced and Expert stages. This behavioural asymptote represents a classic ceiling effect in motor learning theory (Silva, et al., 2020), wherein the physical limits of musculoskeletal execution and sensorimotor conduction come to dominate the timeline. At this advanced stage of

mastery, the cognitive processing speed represents the time required for the ACC to resolve the abstract rule and for M2 to select the corresponding action matrix. This speed likely outpaces the physical mechanics required to orient the body and execute the manual nose-poke response.

Once the physical limits of performance are reached, the brain continues to recruit cellular ensembles (increased Density). This is not done to make the movement faster, but to ensure its stability. By spreading the computational representations across a broader population manifold, the Expert circuit creates high operational dynamic neural recruitment. This backup system shields the behavioural routine against stochastic neural noise, cognitive fatigue, and unexpected environmental distractions. This ensures flawless executive execution, even when behavioural performance metrics appear to have reached a ceiling.

4.3 From perseveration to anticipation: Error analysis

The evolution of error types provides a window into the rat's changing mental state. In our analyses, rather than treating errors as a homogenous metric of performance failure, we segregated mistakes into perseverative and regressive categories. This method of analyses exposed a shift in how the ACC-M2 circuit handles retroactive interference and environmental volatility.

The suppression of perseverative interference

Perseverative errors are characterized by the persistent execution of a historically reinforced rule immediately following a structural switch. These errors provide a classical representation of cognitive inflexibility. In our task, perseverative errors decreased sharply across stages. In the Beginner phase, the prevalence of these errors reflects a failure of top-down

inhibitory gating; the animal was temporarily incapable of suppressing the behavioural momentum of the previous task set. This specific executive deficit is seen in models of addiction and neuropsychiatric disorders, where individuals fail to update behavioural policies despite altered reward contingencies (Landry, 2021). Across time and Stage, the sharp, non-linear drop in perseverative errors provides direct empirical proof of "learning to learn".

Neurobiologically, this clearing of proactive interference (Guise & Shapiro, 2017) aligns with the architectural division of labour observed in our single-unit data. The stable, upstream single-unit scaffolding of the ACC detects the context change and signals the need for behavioural modification. This allows the highly plastic downstream M2 units to efficiently suppress the outdated motor plan and route actions toward the newly relevant spatial targets (Takeuchi, et al., 2022).

The paradox of Expert regressive errors

Conversely, our analysis unmasked a paradoxical surge in regressive errors during the Expert stage. Regressive errors are mistakes made after an animal has already demonstrated a successful switch. Traditional categorical frameworks typically classify these errors as a sign of cognitive relapse or unstable set maintenance. Within the high-density trial architecture of this switching paradigm, this late-Stage error spike points to a different phenomenon, potentially reflecting Anticipatory Interference (Nicholson, Karayanidis, Poboka, Heathcote, & Michie, 2005; Wadsley, Cirillo, Nieuwenhuys, & Byblow, 2023).

4.4 Rapid communication: Similar population coding of task variables

The primary objective of the dPCA analyses was to characterize how task-relevant information was represented at the population level within ACC and M2 and to identify neurons contributing strongly to these representations. Overall, the results revealed remarkably similar coding structures across the two regions. Although ACC and M2 are often associated with distinct functional roles in cognitive control and action selection, respectively, the present analyses did not reveal a strong population-level dissociation between the regions.

Across both Delay and Response epochs, condition-independent activity accounted for the largest proportion of explained variance. This finding suggests that much of the neural population activity reflected processes shared across task conditions, including temporal progression through the trial, anticipation of upcoming events, and preparation for behavioural responses. Similar condition-independent dynamics have been reported in previous population-level studies and are thought to reflect fundamental aspects of neural computation that are not tied to specific task variables (Churchland, et al., 2012; Kobak, et al., 2016; Vyas, Golub, Sussillo, & Shenoy, 2021).

Among task-dependent factors, Side information consistently accounted for more variance than Rule information in both ACC and M2. During the Delay period, Side-related activity represented the largest task-dependent source of variation, while during the Response period Side coding became even more prominent. This result suggests that neural populations in both regions strongly represented information related to the animal's eventual behavioural choice or movement direction. In contrast, Rule-related activity accounted for a relatively small proportion of explained variance across both regions and epochs. Although Rule-selective components were clearly present, their contribution was substantially smaller than that of Side-related and

condition-independent activity. This observation indicates that Rule information was represented within the neural populations but did not dominate overall population structure. Instead, Rule signals appeared to be embedded within a broader network representation that was influenced by temporal and choice-related activity. Interaction-related components also explained relatively little variance. While Side x Rule interactions were detectable in both regions, these effects contributed modestly to overall population activity. This finding suggests that Side and Rule information were represented largely independently at the population level rather than through strong nonlinear combinations of the two variables.

One of the most notable findings was the degree of similarity observed between the ACC and M2. The relative contributions of condition-independent, Side-related, Rule-related and interaction-related activity was broadly comparable across regions. Furthermore, the temporal evolution of the major demixed components exhibited similar patterns in both the ACC and M2. These results suggest that, within the context of the present rule-switching task, both regions participated in representing overlapping aspects of task information. Importantly, the purpose of the dPCA analyses were exploratory. They were intended to provide a population-level overview of neural activity and to identify neurons that contributed strongly to specific dimensions of population variance. Consequently, the findings should not be interpreted as evidence that ACC and M2 perform identical functions. Rather, they demonstrate that population-level representations of Side and Rule information were organized similarly within the two regions under the conditions examined here.

As aforementioned, the dPCA results also provided a principled framework for selecting exemplar neurons for subsequent analyses. By identifying neurons with the largest loadings on

individual components, it was possible to examine how population-level representations emerged from the activity of individual cells.

4.5 The Relay Race of information: Dynamic neural recruitment as a learning strategy

One striking finding from the auROC analysis is that behavioural expertise is not characterized by neurons firing louder, harder, or more accurately, but by them firing longer and in greater numbers. This is supported by the significant Stage by Epoch interaction ($p = 0.0262$) which shows that the active time windows expand and flatten. Instead of a brief, disjointed burst of activity right before the trial starts, the Expert network front-loads and sustains that information straight through the outcome evaluation window (**Figure *11***). Neurons are maintaining the representation for a longer, continuous duration across the ITI. If the brain learned by sharpening the discriminative power of individual units, we would have observed a significant increase in Mean Strength (auROC) as rats transitioned from Beginner to Expert. Instead, we found that the individual signal quality remains remarkably static in the ACC. This suggests that learning to perform a complex rule-based task is akin to a relay race. In the Early stage of training, the "information baton"—the neural representation of the Rule or Side—is held by a small number of neurons for short, disjointed bursts. These neurons may fire significantly, but they fail to maintain the representation consistently across the ITI (**Figure *11***). By the Late stage, the circuit has recruited a much larger and more diverse ensemble of neurons to participate within the execution network. While each individual runner in this relay is no faster than those in the Early stage, the hand-offs are more frequent, and the coverage is more complete.

This population-level expansion ensures that the task-relevant information is never "dropped", effectively bridging the gap between the consumption of a reward and the preparation for the subsequent trial (Braver, 2012). This transition from sparse to dense representation suggests that the hallmark of a trained brain is temporal recruitment, which provides the robustness necessary to guide behaviour in noisy, real-world environments.

4.6 A unified frontal circuit: Asymmetric network optimization

A secondary but equally critical observation from the auROC analysis is the nuanced spatial and temporal division of labour within the circuit. At a single-unit level, a Stage x Region interaction unmasked targeted execution-level plasticity specific to M2 ($p = 0.0447$), while individual ACC units remained remarkably stable ($p = 0.901$). This indicates that the ACC serves as a steady cognitive manager of task rules, whereas M2 functions as the highly plastic execution site that amplifies its raw tuning over weeks (Barthas & Kwan, 2017). At the population level, while the global Stage main effect confirmed an increase in active interval density across the ITI, post-hoc contrasts revealed that this expansion was strongly asymmetric, driven by cellular recruitment in M2 while the ACC population stayed flat. This regional rebalancing alters the functional network: Beginner animals rely heavily on the ACC during Feedback monitoring, whereas Expert animals transition into an M2-dominated Preparatory framework.

For the sub-group of rats ($n = 5$) where trial-by-trial sampling density was maximally preserved, the presence of a significant three-way interaction (Stage:Region:CompType; $p = 0.0479$) reveals that this manifold expansion is not a generic, unguided increase in baseline activity. Instead, it represents a structured reconfiguration of the network's geometry. The

positive interaction coefficient (β = 3.55%) demonstrates that Late-Stage recruitment of M2 is heavily customized to support egocentric action coordinates (CompType_Side). This aligns well with the known anatomical role of M2 in mapping motor planning pathways (Barthas & Kwan, 2017). Meanwhile, the ACC maintains balanced coverage across abstract Rules and Side. This confirms that as cognitive expertise develops, the ACC-M2 projection space selectively channels specific task variables into their appropriate functional nodes—a geometric routing scheme that supports the group-wide stabilization of inter-trial representations (N = 11) during task mastery.

Thus, the transition from a Beginner to an Expert state is not characterized by a uniform, network-wide intensification of single neurons firing harder. Instead, while M2 selectively sharpens its individual output and expands its population footprint, the broader network optimizes by learning to coordinate and open more parallel processing channels across the collective population manifold. This population expansion creates a stable, distributed, and noise-resistant neural workspace capable of flexible, rule-guided behaviour.

4.7 Limitations

Most limitations of this study relate to generalizability of the data both to the larger rat population and inter-species. Ecological validity is limited because the win-shift/win-stay task may not have fully mimicked the natural foraging behaviour that we see in wild rats. This means the experiment may not have reflected the way that rats use cognitive flexibility in real-world settings. And while rats are commonly used in cognitive flexibility research, the extent to which findings can be generalized to other species, including humans, remains unclear. The

mechanisms underlying cognitive flexibility in rats may differ significantly from those found in primates or humans due to differences in brain structure and function.

There is also the limitation of uncontrollable confounding variables such as motivation levels and stress levels of the rats. Studies have shown that if motivation is low, subjects will have poorer engagement in the cognitive training task (Lee, et al., 2024). To mitigate this limitation, our experiment incorporated various measures to encourage motivation levels to be at their highest. To ensure engagement, we reduced outside distractions to a minimum, with rats inside an operant box which obscured distractions and outside interactions. Previous studies have shown that cognitive training improves performance in sustained attention and concentration, which increases motivation during training (Lee, et al., 2024).

Motivation to learn and continue to train can be enhanced by increasing engagement in cognitive training programs, in turn, enhancing effectiveness (Lee et al., 2024). Hence, we ensured that rats had adequate training by satisfying training goals according to performance and not timelines. Motivation can also be reduced or enhanced based on rewards given on task performance (Lee, et al., 2024). Not offering sufficient rewards reduces motivation as the animal does not see worth in completing the task. Contrarily, offering excessive external reward for a behaviour that has already been internally rewarded may also reduce intrinsic motivation (Lee et al., 2024). We maintained consistent and sufficient reward delivery throughout training.

Task load also relates to motivation, with research suggesting that high task loads yield heightened psychological pressure and thus, decreased cognitive flexibility (Wang, Zhu, Zuo, Liu, & Liu, 2023; Wickens, 2008). This relates to the neuropsychiatric disorder of generalized anxiety, as presence of anxiety disorder also accompanies decreased cognitive flexibility. Hence, the more anxious a participant feels during a given task, the less flexible their cognition will be

(Hyer, et al., 2021). Moreover, the subjects are a limitation as we only experimented with male rats. This means that we cannot draw any conclusions about sex differences that could have influenced performance (e.g., hormonal cycles, stress responses, learning strategies).

Task logistics yield limitations of their own; rats responding with a nose-poke instead of using other response communication strategies such as a lever may have contributed to perseverative errors. Previous research has shown that, overall, rats using the nose-poke response method tend to make more errors than rats using the lever response method, during acquisition (Smith, Pattison, & Zentall, 2016). This is thought to be because of the advantage of body-orienting cues in the lever method, which is not present in the nose-poke method, produces faster learning in similar tasks to the one we used during this experiment (Smith, Pattison, & Zentall, 2016). That said, it is important to add that the results of the experiment by Smith et al. (2016) showed that, although the lever group showed better performance globally and locally prior to reversal, the nose-poke group showed better performance locally following the reversal. This is thought to be because of the response separation between paw versus head (Smith, Pattison, & Zentall, 2016). Thus, although the nose-poke method may enhance the number of errors committed during acquisition, it also has its perks.

The final limitation that we will address here is the potential of habit formation. Over time, rats can develop habits or automatic behaviours in win-shift/win-stay tasks (Clemens, Castino, Cornish, Goodchild, & Holmes, 2014). This may reduce the sensitivity to the task in measuring true cognitive flexibility as, instead of continuously adapting their behaviour, the rats may fall into habitual patterns that no longer reflect flexible cognitive processing.

4.8 Relevance and application

Results seen in the behavioural results of the current experiment translate directly to human research of set-shifting and flexibility, demonstrating that with training, the accuracy of tasks requiring cognitive flexibility will increase and both the reaction time and number of overall, general errors will decrease. These results are helpful for situations where individuals have impaired cognitive flexibility, such as in-patients with mental health difficulties like anorexia nervosa (Sato, et al., 2013).

Furthermore, the population-level expansion from the neural results has clinical implications. The observed improvement in flexibility suggests that the frontal circuit's capacity to maintain a mental set is a trainable skill, directly mirroring the human ability to reduce switch costs through practice. The results of the current research are also helpful for research in the aging of the human brain as, as human-beings age, their ability to flexibly Shift between tasks declines (Grönholm-Nyman, et al., 2017; Uddin, 2021). Introducing applications or real-life tasks that actively improve cognitive flexibility could help promote human set Shifting abilities in the cognitively impaired, while also helping to maintain human set Shifting abilities for longer periods of time (into older age). Past research supports this in showing that computerized set Shifting training in the elderly shows long-lasting effects on the trained tasks (Grönholm-Nyman, et al., 2017). It is hypothesized that these set Shifting abilities can be translated from computerized training over to real-world events because previous research has suggested that task-switching training improves performance not only for the trained tasks, but also for structurally similar or even dissimilar tasks (Jaeggi, Buschkuehl, Jonide, & Perrig, 2008; Kattner, Samaan, & Schubert, 2019). However, this needs to be researched further to ascertain its effect on generalized tasks.

4.9 Future directions

The findings established in this project provide a foundation for how the ACC-M2 circuit optimizes task representation, showing that rule-switching mastery is supported by a structural expansion of population-level informational density (PercSig) that is asymmetric and heavily anchored by M2. However, while our auROC and LME modeling proved that a larger pool of neurons is recruited across training, these methods treat population activity as a collection of independent channels. They cannot capture how these units interact, coordinate, or synchronize in real time. To resolve the network-level architectures driving cognitive flexibility, a critical future direction will be to implement neural ensemble analyses. Rather than analyzing single-unit properties in isolation, future work will focus on tracking the co-firing statistics and collective structural organization of simultaneously recorded neurons within and across the ACC and M2.

In this study, we had a modest unit count (< 50 neurons per region per rat total), thus, our future direction will focus on proportions and timings. This can be divided into four phases. Phase (1) will involve the data preparation and quality control, where we will establish rigorous, pre-registered inclusion thresholds. We will use stability and cell typing to filter out units showing firing rate drift or loss across sessions. We can use the spike duration (trough-to-peak width) to classify units into putative pyramidal cells versus interneurons and generate a standard table reporting total unit counts per region per session for transparency. In phase (2) we will expand the behavioural profile to unpack the exact mechanics of rule-switching efficiency. This involves the expansion of our behavioural analyses to include switch cost analysis. For this, we align trials to the within-session rule switch (trials -5 to +20) to track how many trials it takes the animal to recover to baseline accuracy. We will implement a Win-Stay/Lose-Shift analysis to see

how reward-driven strategies alter with training. We will employ a logistic regression to model choices using previous choices, outcomes, stimuli, and current rules to quantify shifting cognitive weights across days. For phase (3) we will perform core neural analyses. Single-session populations are small and thus we will use pool units across sessions and rats. We will track practice effects by binning training sessions into phases (e.g., Early: days 1-5, Middle: days 6-10, Late: days 11-15). We will test if training shrinks the latency of rule coding in the ACC or accelerates prospective choice coding in M2. Lastly, we will perform phase 4 where we will use switch-aligned reconfiguration. We will define a switch epoch (trials -10 to +20) around the rule change, identify switch-responsive neurons that abruptly ramp or step their firing rates, and determine if ACC shifts its firing immediately at trial +1 (post-error) while M2 updates later, and whether this temporal lag shrinks as the animal becomes an expert.

Moreover, future studies can further investigate error-related behaviour as it pertains to this novel task. In a rule-switching paradigm such as this one, errors are not simply mistakes, instead, they are highly informative data points for the animal. As seen in the current paper's first chapter, we took a brief look into error behaviour, analyzing perseveration and regression. That said, analyzing further error behaviour in this task would offer a clearer window into how the ACC-M2 circuit updates behavioural strategies in real-time. For example, one could compare perseveration to exploration, determining whether errors are occurring due to the animal adhering to the previously rewarded rule or searching for the new rule structure. Tracking the ACC-M2 dynamics during these specific error trials could reveal how evaluative conflict signals in the ACC are translated into behavioural shifts in M2.

When looking at error-related behaviour, it would also be worth examining neural activity during post-error slowing (PES)—where an animal exhibits significantly longer reaction times or

decision-making on the trial immediately following an error compared to following a correct response (Rabbitt, 1966). The PES is not thought to be a sign of motor impairment, but instead it is widely recognized as a cognitive control mechanism (Danielmeier & Ullsperger, 2011). Examining the PES would elucidate the network-level mechanisms underlying the deliberate pauses animals take to recalculate strategy following a mistake. Finally, it would be interesting to investigate at the electrophysiological markers of error processing, specifically error-related negativity (ERN) and error positivity (Pe). The ERN is a rapid, negative deflection thought to be generated primarily in the ACC, which occurs ~100 milliseconds after an incorrect choice (Falkenstein, Hohnsbein, Hoormann, & Blanke, 1991). In our rule-switching paradigm, the magnitude of the ERN could indicate how quickly the rat's ACC detects that a previously successful rule is no longer valid. The Pe follows the ERN and is a slower, positive waveform peaking ~200-500 milliseconds post-error (Falkenstein, Hohnsbein, Hoormann, & Blanke, 1991). While the ERN reflects the initial, subconscious detection of an error, the Pe is widely associated with conscious error awareness (Steinhauser & Yeung, 2010). Within the ACC-M2 circuit, this two-stage waveform sequence provides a precise temporal blueprint of adaptation and the neural basis of performance monitoring (Godlove, et al., 2011).

Future studies should investigate whether improvement in flexibility transfers to other tasks in rats. Despite receiving training and practice in the operant box, would cognitive flexibility improve in an alternate task? Notable evidence has suggested that this transfer is possible in training-related improvement of executive functions to untrained tasks in humans (Kattner, Samaan, & Schubert, 2019) but is it possible for cognitive flexibility; an emergent property of efficient executive function (Dajani & Uddin, 2015). It is worth noting that it is difficult to isolate cognitive flexibility from general executive function as it requires the

confluence of several aspects of executive function (Uddin, 2021). Furthermore, research by Grönholm-Nyman and colleagues (2017) has shown that longitudinally (1-year post-training), computerized set Shifting training in the elderly shows long-lasting effects on the trained tasks but very little benefit in terms of generalization. Does this mean that the improvement in flexibility is transferrable but not long-term? The future study must be longitudinal to deepen existing literature. To carry out these future studies, rats would need to be trained on one task, but tested on another, both immediately after training and intermittently for a minimum of 1-year after training. Lastly, future direction should analyze reaction time spiking to determine the cause of the regressive error increase in later stages. This can rule out the concern that these errors are caused by attention deficit or failure to comprehend the task.

A critical question remains: is this neural recruitment task-specific, or does it represent a fundamental strengthening of executive control? Future research must determine if the expert-level dynamic neural recruitment developed here transfers to novel, dissimilar tasks—a concept known as far transfer. In other words, can targeted training on a singular rule-switch paradigm enhance the general set-shifting abilities of a human being, or is the brain simply becoming more efficient at managing a specific set of tools? Understanding whether this localized increase in informational density is a global or local phenomenon will be the next frontier in the study of cognitive rehabilitation and human performance.

4.10 Conclusion

In conclusion, the rats in this study demonstrated remarkable cognitive flexibility, successfully navigating the complex transitions between Stay and Shift rules. These findings

provide empirical support for the Adaptive Coding Hypothesis (Duncan, 2001), which posits that neurons within the PFC possess a high degree of functional plasticity. Rather than being hard-wired for a single stimulus, these units dynamically reconfigure their response properties to prioritize information relevant to the current task context. Our results extend this theory by demonstrating that this adaptation is a specialized division of labour: while upstream ACC units function as a stable baseline scaffold, learning selectively sharpens individual downstream M2 responses alongside an asymmetric population-level expansion within the motor network to lock in predictive temporal timing during expertise acquisition.

Credit authorship contribution statement

Liya Ma and Yiping Zhang contributed to the study conception and design. Experiments and data collection were performed by Liya Ma and Yiping Zhang. Data analysis was performed by Ryley Nathaniel and Liya Ma. The first draft of the manuscript was written by Ryley Nathaniel, while Liya Ma provided feedback. All authors approved the final manuscript.

Declaration of competing interest

There is no competing interest.

Data availability

Data will be made available on request.

Acknowledgement

This work was supported by the Connected Minds Master's Scholarship. Data collection was performed at the Radboud University Nijmegen located in Nijmegen, Netherlands. Data analysis and manuscript writing was performed at York University located in Toronto, Canada.

References

Alexander, W. H., & Brown, J. W. (2019). The role of the anterior cingulate cortex in prediction error and signaling surprise. *Top Cogn Sci., 11*(1), 119-135. doi:https://doi.org/10.1111/tops.12307

Barlow, H. B. (1961). Possible principles underlying the transformations of sensory messages. In H. B. Barlow, & W. A. Rosenblith, *Sensory Communication* (pp. 217-234). Cambridge, MA: MIT Press. doi:https://doi.org/10.7551/mitpress/9780262518420.003.0013

Barthas, F., & Kwan, A. C. (2017). Secondary motor cortex: Where 'sensory' meets 'motor' in the rodent frontal cortex. *Trends Neurosci, 40*(3), 181-193. doi:https://doi.org/10.1016/j.tins.2016.11.006

Birrell, J. M., & Brown, V. J. (2000). Medial frontal cortex mediates perceptual attentional set Shifting in the rat. *The Journal of Neuroscience, 20*(11), 4320-4324. doi:https://doi.org/10.1523/JNEUROSCI.20-11-04320.2000

Botvinick, M. M., Braver, T. S., Barch, D. M., Carter, C. S., & Cohen, J. D. (2001). Conflict monitoring and cognitive control. *Psychol Rev., 108*(3), 624-652. doi:https://doi.org/10.1037/0033-295x.108.3.624

Brady, A. M., & Floresco, S. B. (2015). Operant procedures for assessing behavioral flexibility in rats. *J Vis Exp, 96*, 52387. doi:https://doi.org/10.3791%2F52387

Braver, T. S. (2012). The variable nature of cognitive control: a dual mechanisms framework. *Trends Cogn Sci., 16*(2), 106-113. doi:https://doi.org/10.1016/j.tics.2011.12.010

Brigman, J. L., Graybeal, C., & Holmes, A. (2010). Predictably irrational: Assaying cognitive inflexibility in mouse models of schizophrenia. *Front Neurosci, 4*(13). doi:https://doi.org/10.3389%2Fneuro.01.013.2010

Cantwell, A., Buckholtz, J. W., Atencia, R., & Rosati, A. G. (2022). The origins of cognitive flexibility in chimpanzees. *Developmental Science, 25*(5), e13266. doi:https://doi.org/10.1111/desc.13266

Carter, C. S., Braver, T. S., Barch, D. M., Botvinick, M. M., Noll, D., & Cohen, J. D. (1998). Anterior cingulate cortex, error detection, and the online monitoring of performance. *Science, 280*(5364), 747-749. doi:https://doi.org/10.1126/science.280.5364.747

Cho, K. K., Davidson, T. J., Bouvier, G., Marshall, J. D., Schnitzer, M. J., & Sohal, V. S. (2020). Cross-hemispheric gamma synchrony between prefrontal parvalbumin interneurons supports behavioral adaptation during rule shift learning. *Nat Neurosci., 23*(7), 892-902. doi:https://doi.org/10.1038/s41593-020-0647-1

Chugani, H. T. (1999). Review: Metabolic imaging: A window on brain development and plasticity. *The Neuroscientist, 5*(1), 29–40. doi:https://doi.org/10.1177/107385849900500105

Churchland, M. M., Cunningham, J. P., Kaufman, M. T., Foster, J. D., Nuyujukian, P., Ryu, S. I., & Shenoy, K. V. (2012). Neural population dynamics during reaching. *Nature, 487*, 51–56. doi:https://doi.org/10.1038/nature11129

Clemens, K. J., Castino, M. R., Cornish, J. L., Goodchild, A. K., & Holmes, N. M. (2014). Behavioral and neural substrates of habit formation in rats intravenously self-administering nicotine. *Neuropsychopharmacology, 39*(11), 2584-2593. doi:https://doi.org/10.1038%2Fnpp.2014.111

Cochereau, J., Zandvoort, M., Santariius, T., & Mandonnet, E. (2021). Set Shifting. *Intraoperative Mapping of Cognitive Networks*, 273-285. doi:https://doi.org/10.1007/978-3-030-75071-8_17

Constantinidis, C., & Klingberg, T. (2016). The neuroscience of working memory capacity and training. *Nature Reviews Neuroscience, 17*, 438–449. doi:https://doi.org/10.1038/nrn.2016.43

Cunningham, J. P., & Yu, B. M. (2014). Dimensionality reduction for large-scale neural recordings. *Nat Neurosci., 17*(11), 1500–1509. doi:https://doi.org/10.1038/nn.3776

Dajani, D. R., & Uddin, L. Q. (2015). Demystifying cognitive flexibility: Implications for clinical and developmental neuroscience. *Trends Neurosci., 38*(9), 571-578. doi:https://doi.org/10.1016%2Fj.tins.2015.07.003

Danielmeier, C., & Ullsperger, M. (2011). Post-error adjustments. *Front Psychol., 2*, 233. doi:https://doi.org/10.3389/fpsyg.2011.00233

Dias, R., Robbins, T. W., & Roberts, A. C. (1996). Dissociation in prefrontal cortex of affective and attentional shifts. *Nature, 380*, 69-72. doi:https://doi.org/10.1038/380069a0

Dias, R., Robbins, T. W., & Roberts, A. C. (1997). Dissociable forms of inhibitory control within prefrontal cortex with an analog of the Wisconsin Card Sort test: Restriction to novel situations and independence from “on-line” processing. *The Journal of Neuroscience, 17*(23), 9285-9297. doi:https://doi.org/10.1523/JNEUROSCI.17-23-09285.1997

Ding, M., Tomsick, P. L., Young, R. A., & Jadhav, S. P. (2025). Ventral tegmental area dopamine neural activity switches simultaneously with rule representations in the medial prefrontal cortex and hippocampus. *J Neurosci.*, e1670242025. doi:https://doi.org/10.1523/jneurosci.1670-24.2025

Donoghue, J. P., & Wise, S. P. (1982). The motor cortex of the rat: Cytoarchitecture and microstimulation mapping. *J Comp Neurol., 212*(1), 76-88. doi:https://doi.org/10.1002/cne.902120106

Driscoll, L. N., Pettit, N. L., Minderer, M., Chettih, S. N., & Harvey, C. D. (2017). Dynamic reorganization of neuronal activity patterns in parietal cortex. *Cell., 170*(5), 986-999. doi:https://doi.org/10.1016/j.cell.2017.07.021

Duan, C. A., Erlich, J. C., & Brody, C. D. (2015). Requirement of prefrontal and midbrain regions for rapid executive control of behavior in the rat. *Neuron, 86*(6), 1491-1503. doi:https://doi.org/10.1016/j.neuron.2015.05.042

Duan, C. A., Pagan, M., Piet, A. T., Kopec, C. D., Akrami, A., Riordan, A. J., . . . Brody, C. D. (2021). Collicular circuits for flexible sensorimotor routing. *Nature Neuroscience, 24*, 1110-1120. doi:https://doi.org/10.1038/s41593-021-00865-x

Duncan, J. (2001). An adaptive coding model of neural function in prefrontal cortex. *Reviews Neuroscience, 2*, 820-829. doi:https://doi.org/10.1038/35097575

Duncan, J. (2010). The multiple-demand (MD) system of the primate brain: mental programs for intelligent behaviour. *Trends Cogn Sci, 14*(4), 172-179. doi:https://doi.org/10.1016/j.tics.2010.01.004

Durstewitz, D., Vittoz, N. M., Floresco, S. B., & Seamans, J. K. (2010). Abrupt transitions between prefrontal neural ensemble states accompany behavioral transitions during rule learning. *Neuron, 66*(3), 438-448. doi:https://doi.org/10.1016/j.neuron.2010.03.029

Erlich, J. C., Bialek, M., & Brody, C. D. (2011). A cortical substrate for memory-guided orienting in the rat. *Neuron., 72*(2), 330-343. doi:https://doi.org/10.1016/j.neuron.2011.07.010

Erlich, J. C., Brunton, B. W., Duan, C. A., Hanks, T. D., & Brody, C. D. (2015). Distinct effects of prefrontal and parietal cortex inactivations on an accumulation of evidence task in the rat. *eLife, 4*, e05457. doi:https://doi.org/10.7554/eLife.05457

Falkenstein, M., Hohnsbein, J., Hoormann, J., & Blanke, L. (1991). Effects of crossmodal divided attention on late ERP components: II. Error processing in choice reaction tasks. *Electroencephalogr Clin Neurophysiol, 78*(6), 447–455. doi:https://doi.org/10.1016/0013-4694(91)90062-9

Floresco, S. B., Block, A. E., & Tse, M. T. (2008). Inactivation of the medial prefrontal cortex of the rat impairs strategy set-shifting, but not reversal learning, using a novel, automated procedure. *Behav. Brain Res., 190*(1), 85-96. doi:https://doi.org/10.1016/j.bbr.2008.02.008

Friedman, N. P., & Robbins, T. W. (2022). The role of prefrontal cortex in cognitive control and executive function. *Neuropsychopharmacology, 47*, 72–89. doi:https://doi.org/10.1038/s41386-021-01132-0

Ghods-Sharifi, S., Haluk, D. M., & Floresco, S. B. (2008). Differential effects of inactivation of the orbitofrontal cortex on strategy set-shifting and reversal learning. *Neurobiol Learn Mem., 89*(4), 567-573. doi:https://doi.org/10.1016/j.nlm.2007.10.007

Godlove, D. C., Emeric, E. E., Segovis, C. M., Young, M. S., Schall, J. D., & Woodman, G. F. (2011). Event-related potentials elicited by errors during the Stop-Signal Task. I. Macaque monkeys. *J Neurosci., 31*(44), 15640–15649. doi:https://doi.org/10.1523/JNEUROSCI.3349-11.2011

Grant, D. A., & Berg, E. (1948). A behavioral analysis of degree of reinforcement and ease of shifting to new responses in a Weigl-type card-sorting problem. *Journal of Experimental Psychology, 38*(4), 404–411. doi:https://psycnet.apa.org/doi/10.1037/h0059831

Grönholm-Nyman, P., Soveri, A., Rinne, J. O., Ek, E., Nyholm, A., Neely, A. S., & Laine, M. (2017). Limited effects of set Shifting training in healthy older adults. *Front. Aging Neurosci., 9*. doi:https://doi.org/10.3389%2Ffnagi.2017.00069

Guise, K. G., & Shapiro, M. L. (2017). Medial prefrontal cortex reduces memory interference by modifying hippocampal encoding. *Neuron., 94*(1), 183-192. doi:https://doi.org/10.1016/j.neuron.2017.03.011

Hamilton, D. A., & Brigman, J. L. (2016). Behavioral flexibility in rats and mice: Contributions of distinct frontocortical regions. *Genes Brain Behav., 14*(1), 4-21. doi:https://doi.org/10.1111%2Fgbb.12191

Huttenlocher, P. R. (1984). Synapse elimination and plasticity in developing human cerebral cortex. *American Journal of Mental Deficiency, 88*(5), 488–496. Retrieved from https://pubmed.ncbi.nlm.nih.gov/6731486/

Hyer, M. M., Shaw, G. A., Goswamee, P., Dyer, S. K., Burns, C. M., Soriano, E., . . . Neigh, G. N. (2021). Chronic adolescent stress causes sustained impairment of cognitive flexibility and hippocampal synaptic strength in female rats. *Neurobiol Stress., 14*, 100303. doi:https://doi.org/10.1016/j.ynstr.2021.100303

Izquierdo, A., Brigman, J. L., Radke, A. K., Rudebeck, P. H., & Holmes, A. (2017). The neural basis of reversal learning: An updated perspective. *Neuroscience, 345*, 12-26. doi:https://doi.org/10.1016/j.neuroscience.2016.03.021

Jaeggi, S. M., Buschkuehl, M., Jonide, J., & Perrig, W. J. (2008). Improving fluid intelligence with training on working memory. *Psychological and Cognitive Sciences, 105*(19), 6829-6833. doi:https://doi.org/10.1073/pnas.0801268105

Kamigaki, T., Fukushima, T., & Miyashita, Y. (2009). Cognitive set reconfiguration signaled by macaque posterior parietal neurons. *Neuron, 61*(6), 941-951. doi:https://doi.org/10.1016/j.neuron.2009.01.028

Kamigaki, T., Fukushima, T., & Miyashita, Y. (2011). Neuronal signal dynamics during preparation and execution for behavioral shifting in macaque posterior parietal cortex. *J Cogn Neurosci., 23*(9), 2503-2520. doi:https://doi.org/10.1162/jocn.2011.21613

Kamigaki, T., Fukushima, T., Tamura, K., & Miyashita, Y. (2012). Neurodynamics of cognitive set shifting in monkey frontal cortex and its causal impact on behavioral flexibility. *J Cogn Neurosci., 24*(11), 2171-2185. doi:https://doi.org/10.1162/jocn_a_00277

Kattner, F., Samaan, L., & Schubert, T. (2019). Cross-modal transfer after auditory task-switching training. *Memory & Cognition, 47*, 1044-1061. doi:https://doi.org/10.3758/s13421-019-00911-x

Kobak, D., Brendel, W., Constantinidis, C., Feierstein, C. E., Kepecs, A., Mainen, Z. F., . . . Machens, C. K. (2016). Demixed principal component analysis of neural population data. *eLife, 5*, e10989. doi:https://doi.org/10.7554/eLife.10989

Kolling, N., Behrens, T., Wittmann, M. K., & Rushworth, M. (2016). Multiple signals in anterior cingulate cortex. *Curr Opin Neurobiol., 37*, 36-43. doi:https://doi.org/10.1016/j.conb.2015.12.007

Konishi, S., Kawazu, M., Uchida, I., Kikyo, H., Asakura, I., & Miyashita, Y. (1999). Contribution of working memory to transient activation in human inferior prefrontal cortex during performance of the Wisconsin Card Sorting Test. *Cerebral Cortex, 9*(7), 745-753. doi:https://doi.org/10.1093/cercor/9.7.745

Kramer, J. H., Quitania, L., Dean, D., Neuhaus, J., Rosen, H. J., Halabi, C., . . . Miller, B. L. (2008). Magnetic resonance imaging correlates of set shifting. *J Int Neuropsychol Soc., 13*(3), 386-392. doi:https://doi.org/10.1017%2FS1355617707070567

Landry, O. (2021). An examination of perseverative errors and cognitive flexibility in autism. *PLoS One, 16*(1), e0223160. doi:https://doi.org/10.1371%2Fjournal.pone.0223160

Laubach, M., Amarante, L. M., Swanson, K., & White, S. R. (2018). What, if anything, is rodent prefrontal cortex? *eNeuro, 5*(5), ENEURO.0315-18.2018. doi:https://doi.org/10.1523/ENEURO.0315-18.2018

Lee, L. Y., Healy, M. P., Fischer, N. L., Tong, K., Chen, A. S., Sahakian, B. J., & Kourtzi, Z. (2024). Cognitive flexibility training for impact in real-world settings. *Current Opinion in Behavioral Sciences, 59*, 101413. doi:https://doi.org/10.1016/j.cobeha.2024.101413

Ma, L., Hyman, J. M., Durstewitz, D., Phillips, A. G., & Seamans, J. K. (2016). A quantitative analysis of context-dependent remapping of medial frontal cortex neurons and ensembles. *Journal of Neuroscience, 36*(31), 8258-8272. doi:https://www.jneurosci.org/content/36/31/8258

Mansouri, F. A., Buckley, M. J., Fehring, D. J., & Tanaka, K. (2020). The role of primate prefrontal cortex in bias and shift between visual dimensions. *Cereb Cortex., 30*(1), 85-99. doi:https://doi.org/10.1093/cercor/bhz072

Miles, S., Howlett, C. A., Berryman, C., Nedeljkovic, M., Moseley, G. L., & Phillipou, A. (2021). Considerations for using the Wisconsin Card Sorting Test to assess cognitive

flexibility. *Behavior Research Methods, 53*, 2083-2091. doi:https://doi.org/10.3758/s13428-021-01551-3

Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annu Rev Neurosci., 24*, 167-202. doi:https://doi.org/10.1146/annurev.neuro.24.1.167

Miller, H. L., Ragozzino, M. E., Cook, E. H., Sweeney, J. A., & Mosconi, M. W. (2015). Cognitive set shifting deficits and their relationship to repetitive behaviors in autism spectrum disorder. *J Autism Dev Disord., 45*(3), 805-815. doi:https://doi.org/10.1007/s10803-014-2244-1

Milner, B. (1963). Effects of different brain lesions on card sorting: The role of the frontal lobes. *Arch Neurol., 9*(1), 90-100. doi:10.1001/archneur.1963.00460070100010

Milton, L. K., Mirabella, P. N., Greaves, E., Spanswick, D. C., van den Buuse, M., Oldfield, B. J., & Foldi, C. J. (2021). Suppression of corticostriatal circuit activity improves cognitive flexibility and prevents body weight loss in activity-based anorexia in rats. *Biological Psychiatry, 12*, 819-828. doi:https://doi-org.ezproxy.library.yorku.ca/10.1016/j.biopsych.2020.06.022

Monchi, O., Petrides, M., Petre, V., Worsley, K., & Dagher, A. (2001). Wisconsin card sorting revisited: Distinct neural circuits participating in different stages of the task identified by event-related functional magnetic resonance imaging. *J Neurosci., 21*(19), 7733-7741. doi:https://doi.org/10.1523/JNEUROSCI.21-19-07733.2001

Monsell, S. (2003). Task switching. *Trends Cogn Sci, 7*(3), 134-140. doi:https://doi.org/10.1016/s1364-6613(03)00028-7

Murphy, F. C., Michael, A., & Sahakian, B. J. (2012). Emotion modulates cognitive flexibility in patients with major depression. *Psychol Med., 42*(7), 1373-1382. doi:https://doi.org/10.1017/s0033291711002418

Murphy, F. C., Sahakian, B. J., Rubinsztein, J. S., Michael, A., Rogers, R. D., Robbins, T. W., & Paykel, E. S. (1999). Emotional bias and inhibitory control processes in mania and depression. *Psychological Medicine, 29*, 1307–1321. doi:https://doi.org/10.1017/s0033291799001233

Nakahara, K., Hayashi, T., Konishi, S., & Miyashita, Y. (2002). Functional MRI of macaque monkeys performing a cognitive set-shifting task. *Science, 295*(5559), 1532-1536. doi:https://doi.org/10.1126/science.1067653

Nelson, H. E. (1976). A modified card sorting test sensitive to frontal lobe defects. *Cortex, 12*(4), 313-324. doi:https://www.sciencedirect.com/science/article/pii/S0010945276800354

Nicholson, R., Karayanidis, F., Poboka, D., Heathcote, A., & Michie, P. T. (2005). Electrophysiological correlates of anticipatory task-switching processes. *Psychophysiology, 42*(5), 540-54. doi:https://doi.org/10.1111/j.1469-8986.2005.00350.x

Owen, A. M., Roberts, A. C., Hodges, J. R., & Robbins, T. W. (1993). Contrasting mechanisms of impaired attentional set-shifting in patients with frontal lobe damage or Parkinson's disease. *Brain, 116*(5), 1159-1175. doi:https://doi.org/10.1093/brain/116.5.1159

Owen, A. M., Roberts, A. C., Polkey, C. E., Sahakian, B. J., & Robbins, T. W. (1991). Extra-dimensional versus intra-dimensional set shifting performance following frontal lobe excisions, temporal lobe excisions or amygdalo-hippocampectomy in man. *Neuropsychologia, 29*(10), 993-1006. doi:https://doi.org/10.1016/0028-3932(91)90063-E

Pachitariu, M., Sridhar, S., Pennington, J., & Stringer, C. (2024). Spike sorting with Kilosort4. *Nature Methods, 21*, 914–921. doi:https://doi.org/10.1038/s41592-024-02232-7

Park, J., & Moghaddam, B. (2017). Impact of anxiety on prefrontal cortex encoding of cognitive flexibility. *Neuroscience, 345*, 193-202. doi:https://doi.org/10.1016/j.neuroscience.2016.06.013

Perpiñá, C., Segura, M., & Sánchez-Reales, S. (2017). Cognitive flexibility and decision-making in eating disorders and obesity. *Eating and Weight Disorders - Studies on Anorexia, Bulimia and Obesity, 22*, 435-444. doi:https://doi-org.ezproxy.library.yorku.ca/10.1007/s40519-016-0331-3

Qi, X.-L., & Constantinidis, C. (2013). Neural changes after training to perform cognitive tasks. *Behav Brain Res., 241*, 235–243. doi:https://doi.org/10.1016/j.bbr.2012.12.017

Rabbitt, P. M. (1966). Errors and error correction in choice-response tasks. *Journal of Experimental Psychology, 71*(2), 264–272. doi:https://psycnet.apa.org/doi/10.1037/h0022853

Rady, A., Elsheshai, A., el Wafa, H. A., & Elkholy, O. (2012). WCST performance in schizophrenia and severe depression with psychotic features. *ISRN Psychiatry, 1*, 373748. doi:https://doi.org/10.5402/2012/373748

Rajalakshmi, R., & Jeeves, M. A. (1965). The relative difficulty of reversal learning (reversal index) as a basis of behavioural comparisons. *Animal Behaviour, 13*(2-3), 203-211. doi:https://doi.org/10.1016/0003-3472(65)90035-7

Reindl, E., Voelter, C., Civelek, Z., Duncan, L. F., Lugosi, Z., Felsche, E., . . . Seed, A. M. (2023). The shifting shelf task: a new, non-verbal measure for attentional set shifting. *Proceedings of the Royal Society of London Series B: Biological Sciences, 290*(1991), 20221496. doi:https://doi.org/10.1098/rspb.2022.1496

Rich, E. L., & Shapiro, M. (2009). Rat prefrontal cortical neurons selectively code strategy switches. *J Neurosci., 29*(22), 7208-7219. doi:https://doi.org/10.1523/jneurosci.6068-08.2009

Rich, E. L., & Shapiro, M. L. (2007). Prelimbic/infralimbic inactivation impairs memory for multiple task switches, but not flexible selection of familiar tasks. *J Neurosci., 27*(17), 4747-4755. doi:https://doi.org/10.1523/jneurosci.0369-07.2007

Rigotti, M., Barak, O., Warden, M. R., Wang, X.-J., Daw, N. D., Miller, E. K., & Fusi, S. (2013). The importance of mixed selectivity in complex cognitive tasks. *Nature, 497*(7451), 585-590. doi:https://doi.org/10.1038/nature12160

Rogers, R. D., & Monsell, S. (1995). Costs of a predictible switch between simple cognitive tasks. *Journal of Experimental Psychology: General, 124*(2), 207–231. doi:https://psycnet.apa.org/doi/10.1037/0096-3445.124.2.207

Sato, Y., Saito, N., Utsumi, A., Aizawa, E., Shoji, T., Izumiyama, M., . . . Fukudo, S. (2013). Neural basis of impaired cognitive flexibility in patients with anorexia nervosa. *PLoS One, 8*(5), e61108. doi:https://doi.rg/10.1371/journal.pone.0061108

Shibata, K., Sasaki, Y., Won Bang, J., Walsh, E. G., Machizawa, M. G., Tamaki, M., . . . Watanabe, T. (2017). Overlearning hyper-stabilizes a skill by rapidly making neurochemical processing inhibitory-dominant. *Nat Neurosci., 20*(3), 470-475. doi:https://doi.org/10.1038/nn.4490

Silva, P. R., Silva, R. H., Lima, R. H., Meurer, Y. S., Ceppi, B., & Yamamoto, M. E. (2020). Are there multiple motivators for helping behavior in rats? *Frontiers in Psychology, 11*. doi:https://doi.org/10.3389/fpsyg.2020.01795

Smith, A. P., Pattison, K. F., & Zentall, T. R. (2016). Rats' midsession reversal performance: The nature of the response. *Learning & Behavior, 44*, 49-58. doi:https://doi.org/10.3758/s13420-015-0189-7

St. Onge, J. R., & Floresco, S. B. (2010). Prefrontal cortical contribution to risk-based decision making. *Cerebral Cortex, 20*(8), 1816-1828. doi:https://doi.org/10.1093/cercor/bhp250

Steinhauser, M., & Yeung, N. (2010). Decision processes in human performance monitoring. *Journal of Neuroscience, 30*(46), 15643-15653. doi:https://doi.org/10.1523/JNEUROSCI.1899-10.2010

Stuss, D. T., Levine, B., Alexander, M. P., Hong, J., Palumbo, C., Hamer, L., . . . Izukawa, D. (2000). Wisconsin card sorting test performance in patients with focal frontal and posterior brain damage: Effects of lesion location and test structure on separable cognitive processes. *Neuropsychologia, 38*(4), 388-402. doi:https://doi.org/10.1016/S0028-3932(99)00093-7

Takeuchi, D., Roy, D., Muralidhar, S., Kawai, T., Bari, A., Lovett, C., . . . Tonegawa, S. (2022). Cingulate-motor circuits update rule representations for sequential choice decisions. *Nature Communications, 13*, 4545. doi:https://doi.org/10.1038/s41467-022-32142-1

Tang, H., Riley, M. R., Singh, B., Qi, X.-L., Blake, D. T., & Constantinidis, C. (2022). Prefrontal cortical plasticity during learning of cognitive tasks. *Nature Communications, 13*(90). doi:https://doi.org/10.1038/s41467-021-27695-6

Uddin, L. Q. (2021). Cognitive and behavioural flexibility: neural mechanisms and clinical considerations. *Nature Reviews Neuroscience, 22*, 167-179. doi:https://doi.org/10.1038/s41583-021-00428-w

Uylings, H. B., Groenewegen, H. J., & Kolb, B. (2003). Do rats have a prefrontal cortex? *Behav Brain Res., 146*(1-2), 3-17. doi:https://doi.org/10.1016/j.bbr.2003.09.028

Vallesi, A., Arbula, S., Capizzi, M., Causin, F., & D'Avella, D. (2015). Domain-independent neural underpinning of task-switching: An fMRI investigation. *Cortex, 65*, 173-183. doi:https://doi.org/10.1016/j.cortex.2015.01.016

Velten, J., Gao, X., Van Nierop, P., Sanchez, Y., Domsch, K., Agarwal, R., . . . Lohmann, I. (2022). Single-cell RNA sequencing of motoneurons identifies regulators of synaptic wiring in Drosophila embryos. *Mol Syst Biol., 18*(3), e10255. doi:https://doi.org/10.15252/msb.202110255

Voigts, J., Newman, J. P., Wilson, M. A., & Harnett, M. T. (2020). An easy-to-assemble, robust, and lightweight drive implant for chronic tetrode recordings in freely moving the rats. *Journal of Neural Engineering, 17*(2). doi:https://doi.org/10.1088/1741-2552/ab77f9

Vyas, S., Golub, M. D., Sussillo, D., & Shenoy, K. V. (2021). Computation through neural population dynamics. *Annu Rev Neurosci., 43*, 249–275. doi:https://doi.org/10.1146/annurev-neuro-092619-094115

Wadsley, C. G., Cirillo, J., Nieuwenhuys, A., & Byblow, W. D. (2023). Comparing anticipatory and stop-signal response inhibition with a novel, open-source selective stopping toolbox. *Exp Brain Res., 241*(2), 601–613. doi:https://doi.org/10.1007/s00221-022-06539-9

Wang, Y., Hu, X., & Li, Y. (2022). Investigating cognitive flexibility deficit in schizophrenia using task-based whole-brain functional connectivity. *Frontiers in Psychiatry, 13*. doi:https://doi.org/10.3389/fpsyt.2022.1069036

Wang, Y., Zhu, C., Zuo, D., Liu, J., & Liu, D. (2023). The effect of emotional motivation on strategy flexibility: The moderating role of task load. *Front Psychol., 14*, 1241131. doi:https://doi.org/10.3389/fpsyg.2023.1241131

Wickens, C. (2008). Multiple resources and mental workload. *Human Factors: The Journal of Human Factors and Ergonomic Society, 50*(3), 449-455. doi:https://doi-org.ezproxy.library.yorku.ca/10.1518/001872008X288394

Yu, Y., Yu, Y., & Lin, Y. (2020). Anxiety and depression aggravate impulsiveness: The mediating and moderating role of cognitive flexibility. *Psychology, Health & Medicine, 25*(1), 25-36. doi:https://doi-org.ezproxy.library.yorku.ca/10.1080/13548506.2019.1601748

Zajic, A. (2022). What is Akaike Information Criterion (AIC)? *Built-in*. Retrieved from https://builtin.com/data-science/what-is-aic

Zastrow, A., Kaiser, S., Stippich, C., Walther, S., Herzog, W., Tchanturia, K., . . . Friederich, H.-C. (2009). Neural correlates of impaired cognitive-behavioural flexibility in anorexia

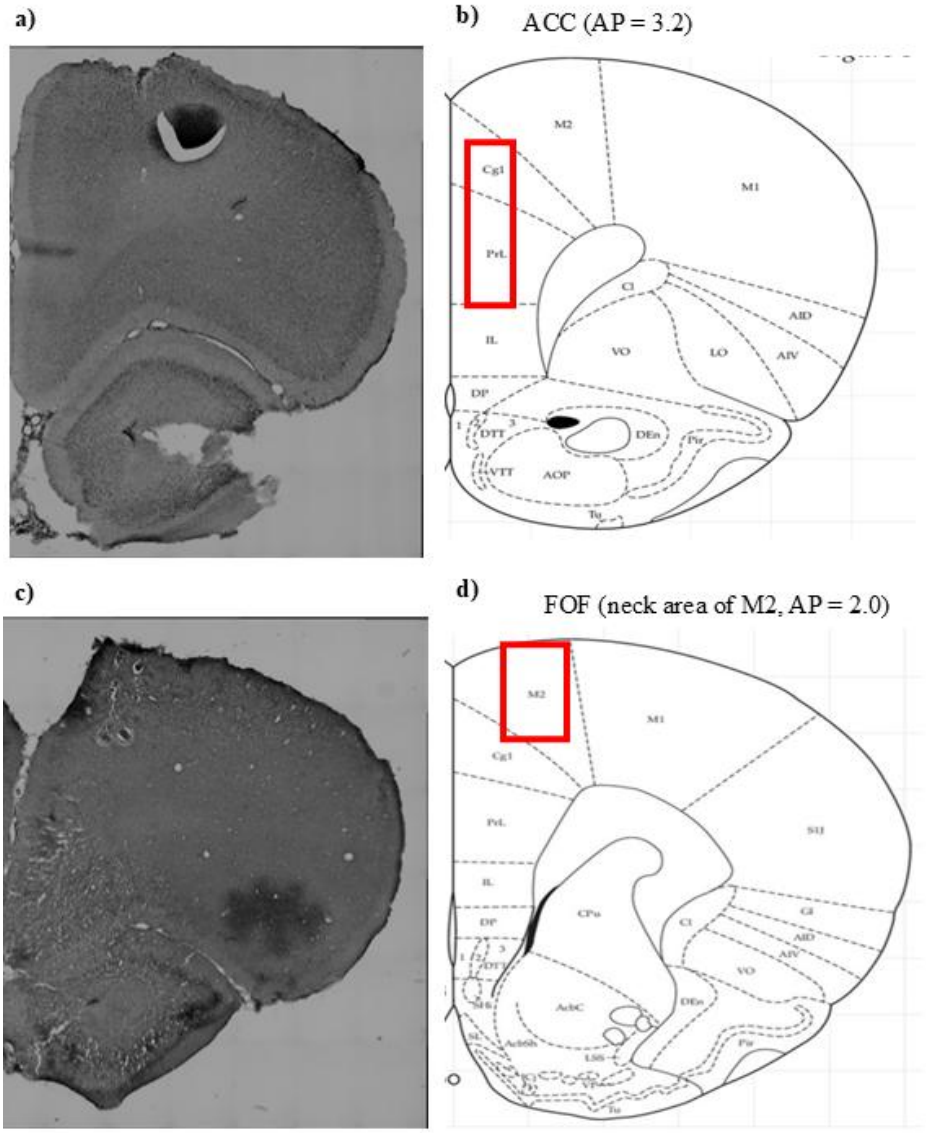
nervosa. *The American Journal of Psychiatry*. doi:https://doi-org.ezproxy.library.yorku.ca/10.1176/appi.ajp.2008.08050775

Zylberberg, J., Pouget, A., Latham, P. E., & Shea-Brown, E. (2017). Robust information propagation through noisy neural circuits. *PLoS Computational Biology, 13*(4), e1005497. doi:https://doi.org/10.1371/journal.pcbi.1005497

Integrated supplementary information

Supplementary figure 1

Histology showing recording sites.



a) Histology showing representative electrode track endings (white arrows) in the ACC. **b**) Schematic showing the location and range of the recording sites in the ACC. Given that not all electrodes left a track, the recording range was inferred based on the visible tracks and the size of the electrode arrays. **c**)

Representative electrode track endings (white arrows) in the M2. **d**) Schematics showing the location and range of the recording sites in the M2.