DISTRACTED AGENTS IN CROWD SIMULATIONS

MELISSA KREMER

A THESIS SUBMITTED TO THE FACULTY OF GRADUATE STUDIES
IN PARTIAL FULFILMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
MASTER OF SCIENCE

GRADUATE PROGRAM IN COMPUTER SCIENCE
YORK UNIVERSITY
TORONTO, ONTARIO

OCTOBER 2020

Abstract

Crowd simulations can provide meaningful insights about crowd movement and flow in arbitrary environments, which is useful for predictive planning and analysis. High fidelity agents and simulations are important factors in the quality of the analysis and their validity. Historically, agents in crowd simulations have been largely homogeneous in terms of physical and behavioural properties. However, this homogenous structure does not accurately reflect crowds seen in the real world, which are largely heterogeneous in a wide variety of aspects including behaviours. In particular, distracted behaviour in crowds has been largely ignored in crowd simulations to date. The use of cellular phones has become one of the most commonly observed distracted behaviour in real world urban crowds. This thesis presents a first step towards modelling this important phenomenon in crowd simulations. The models proposed include a fuzzy goal model for modelling lateral deviation from a straight path reported in distraction literature, and a visual attention and reaction model representing the field of view for distracted and undistracted pedestrians which controls probabilistically when agents stop being distracted and pay attention. Steering parameters are determined by extracting values from real-world distraction studies and thus reflect ground truth. Experiments conducted show that even a few distracted pedestrians exhibiting a handful of distracted activities have a significant impact on several crowd statistics including flow rate, effort, and kinetic energy. Each parameter in this highly parameterizable model is shown to have a significant impact on the resulting simulation and can thus be used to create a variety of resulting crowd simulations.

Acknowledgements

I would first to express my sincerest appreciation and gratitude to my supervisor, Professor Petros Faloutsos. When I look back at when I started as a graduate student at York University, I can scarcely believe how much I've changed in such a short span of time. With his guidance and support I learned and grew so much. He gave me the opportunity to contribute to exciting research in this field, which has led me on a very fulfilling journey and changed the direction of my life.

My deepest appreciation goes to my colleague, mentor, and friend Professor Brandon Haworth at the University of Victoria. His dedication and support empowered me to reach new heights and to bring this thesis to fruition.

I would also like to thank Professor Melanie Baljko for joining my supervisory committee and supporting me throughout the process, as well as my colleagues in the Department of Electrical Engineering and Computer Science at York University. The successful completion of this thesis is due in part by the direction, feedback, and insight they provided me.

Finally, I would like to thank the Department of Electrical Engineering and Computer Science at York University for providing the financial support needed to complete this thesis work.

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

One of the most valuable tools for understanding how people move through a space is crowd simulation. Urban design and evacuation planning benefits from being able to test a design and iterate on it, making modifications and improvements. However, it is often infeasible to validate an environment with a real-world crowd due to costs and safety concerns. By employing crowd simulation techniques, we can mitigate these concerns by testing the environment virtually with simulated crowds.

Various crowd simulation algorithms are used to calculate the movement of virtual agents representing people through a virtual environment. These algorithms take as input initial conditions of the agents, their respective goals, properties of the environment and other parameters in order to calculate agent paths. First, a global path from the agent to the goal is calculated and subsequently divided into multiple sub-goals at every point where the agent should change direction. Then the agent performs local steering to navigate to each sub-goal in turn while performing collision avoidance maneuvers if any other agents or obstacles come into its path. The agent may also be animated as it makes its way toward the goal.

This process is split up into a hierarchy which includes action selection, path planning and steering, and locomotion. This thesis will focus on specific aspects of agent behavior and how steering is affected. Much previous work has been done on agent behavior which will be further discussed in the next chapter. This includes personality and emotion models, psychology, cultural diversity as well as group behaviors and interactions, all of which have been shown to have a strong impact on the overall simulated crowd flow [74, 11, 16, 2, 69, 48]. Even non-normative gaits modelling aspects of medical disabilities can significantly change various measures of simulated crowds [18]. These and other works highlight the importance of heterogeneity and diversity in crowd simulation. However, certain aspects of attention, distraction, and gaze modelling have been ignored or neglected in the literature. In particular, one key aspect of modern pedestrian behavior has remained largely overlooked thus far.

In major cities including Toronto, the increasingly common use of cellphones by pedestrians in crowded urban environments have become a topic of great concern. It was reported that in Toronto in 2008, 598 of 4522 total pedestrian injuries were directly related to distracted pedestrians. The Ontario Provicial Police subsequently responded with the launch of the ”Heads Up” pedestrian safety awareness campaign [41].

Research on cellphone distraction and its impact on crowds can be used to help inform safety and urban design planning. When a pedestrian looks down at their cellphone, their visual field becomes skewed towards the ground, resulting in a reduction in visual information about their environment. Additionally, the cognitive load of the distraction causes them to miss even information that is within their visual field, resulting in a sort of artificial shrinking of the visual field. Thus they don't make as many avoidance behaviors, or are forced to make an abrupt avoidance maneuver at the last second. Spatiotemporal parameters of their gait, such as gait speed and step width, also change. This can have a compounded impact on the overall crowd.

In addition to urban design, modelling distracted pedestrians can improve animations for film and video games, improving realism and fidelity. The following are key contributions that this work has provided:

- An extensive review of real world distraction literature, in which the provided steering model parameters are grounded.
- Modelling of state interactions between agents and their cellphones through the dual use of parametric and interactive event-driven behaviour trees.
- Modelling of lateral deviation from a straight path, as identified in distraction studies, through the use of a parameterizable ellipsoid fuzzy goal model.
- Modelling visual attention by convolving two Gaussians in polar coordinates, which are then used in the steering and attention refocusing models.

- Modelling refocusing attention based on three intrinsic and extrinsic parameters used to perform Bernoulli trials conditioned on the visual attention model.

Portions of this document have been accepted for publication in [34]. This publication includes content from Sections 2.1, 2.2, 3.1, 4.1, and 4.2, and Chapter 5. Concepts explained in Chapters 4, 5, and 6 were jointly developed between the authors Melissa Kremer, Dr. Brandon Haworth, Dr. Mubbasir Kapadia, and Dr. Petros Faloutsos.

Chapter 1 is an introduction to the subject matter, emphasizing the importance heterogeneity in crowd simulations and its role in producing high fidelity simulations that reflect real-world crowds. The discrepancy between the lack of cellphone distracted pedestrians in crowd simulation research and the increasing safety concerns regarding cellphone distracted pedestrians in major cities is discussed. The presented objectives and contributions of this work are a first step towards addressing this gap in knowledge.

Chapter 2 covers the theoretical background of crowd simulation, behaviour and attention modelling. It emphasizes the lack of cellphone distraction modelling in crowd simulation literature, and discusses how this work differs from previously created models.

Chapter 3 discusses literature pertaining to real-world distraction studies. Theoretical concepts such as dual task cost and cognitive loads are discussed. Various observed phenomena including inattentional blindness and cognitive tunnelling are examined. Finally, this chapter includes a literature review in the form of a table detailing values reported in the literature for spatiotemporal gait parameters pertaining to distracted pedestrians.

Chapter 4 explains various aspects of the proposed distraction model. These include a fuzzy goal model based on the real-world observed phenomena of lateral deviation, a visual attention model for distracted agents, and a refocusing model that determines when agents stop being distracted and look up. Additionally, the distraction types included in the experiments with their respective steering parameters are discussed. A crowd authoring tool for generating crowds with distracted agents is also examined.

Chapter 5 examines a set of preliminary experiments that did not include a visual attention model. The experiment design and results are presented and key findings are discussed. The results show the impact of including distracted pedestrians on several crowd flow statistics.

Chapter 6 presents experiments conducted with a visual attention model as an extension to the preliminary experiments. Results demonstrate the various settable distraction parameters and their impact on crowd measures. Key highlights are discussed.

Chapter 7 is the conclusion and summarizes the key points of this work. Possible future work is discussed.

Chapter 2

Related work

Related work for this project includes research in crowd and behaviour modelling, as well as studies on the impact of distraction on individual and crowd behaviours and navigation. Literature on crowd simulation and behaviour modelling in general, as well as gaze and attention modelling specifically has been consulted.

2.1 Crowd Simulation

Over the last few decades, many approaches and techniques have been developed for simulating crowds. It began primarily with Boids [56], an early model that used simple rules to simulate agent navigation and collision avoidance while also being able to maintain flock cohesion behaviour. Force-based models like Social Forces [19] use attractive and repulsive forces to enforce comfortable distances between agents and between agents and obstacles while navigating towards a goal. Anticipatory and synthetic vision based models have also been developed [45, 46, 67, 5], as well as hybrid models [60]. Predictive force-based models were developed as extensions to traditional force-based models. These include the Predictive Avoidance Model (PAM) [30], which in addition to using attractive and repulsive forces to guide agents also uses current trajectories to predict collision-free trajectories for a future time window. PAM is capable of emergent behaviour such as lane formation. Another study improved the flexibility of predictions and the number of factors taken into account by using a probabilistic model [68]. Many other approaches and derivatives of these methods are discussed in the literature, of which a survey has been collected in [29] and in

[63].

2.2 Behaviour Modelling

Incorporating behaviour models into crowd simulations have increased fidelity and accuracy. Cognitive models allow for high-level agent behaviour by using low-level controllers [14]. Decision networks can model interactions between agents with different personalities with corresponding animations [72]. Personality models have also been implemented into crowd simulations and studied for their impact on overall crowd characteristics [12]. Realistic crowd behaviours can also be learned with data-driven methods using real-world crowd data [37, 27]. Many other behaviour models beyond these have been proposed for crowd simulations.

2.3 Gaze and Attention Modelling

Gaze behaviour has been modelled to improve realism in crowd simulations, and is an important consideration for modelling attention and distraction. Some steering algorithms model simplistically model visual field without explicit gaze modelling [35, 50]. Another approach is to use a set of rules to allow agents to shift focus or attention based on external stimuli from the environment [20, 32].

Environment saliency maps have also been used to direct agents' gaze toward the most salient, or most noticeable, object in the field of view [24]. Simulations of face-to-face conversations have incorporated gaze modelling to produce more believable personal interactions [4]. Gaze behaviour has also been implemented to some extent in crowds, or in interactive environments with real users [33, 42]. However, implementations involving distractions in general, and in particular cellphone distractions, have not been investigated to the best of my knowledge in crowd simulations.

Chapter 3

Effects of Distraction in Pedestrian Movement

3.1 Distraction Tasks and Task Costs

It has been shown that performing a secondary task while walking affects gait. Much of the processes involved in walking are subconscious and yet when people attempt to multitask while walking their gait becomes inefficient. Walking and navigating an environment is a complex activity that requires our full attention. When we multitask while walking, our attention is split between the task of walking and the secondary task. We can think of the secondary task in terms of task cost, or how much attention we need to divert away from the primary walking task in order to perform this task in parallel.

In one theory, secondary activities can be visual, motor, or cognitive in nature, or a combination of these. A study found that the type of distracting activity is directly related to the task cost, with cognitive tasks often being the most distracting, followed by motor and then visual [65]. Another study found that gait is affected by cognitive and visual demand of a secondary task involving a cellphone but not affected by gross motor tasks, although no conclusions were drawn about fine motor tasks such as typing and tapping movements of the fingers [54]. Difficulty of the secondary task and its effect on gait has also been examined. Talking on the phone while driving was investigated in [47] which showed that complex conversations had a bigger impact on distraction and reaction time than simple ones. One study found that more difficult secondary tasks had a bigger impact on the gaits of older adults in particular, especially when maneuvering around obstacles [53]. Still more research may be required to fully understand the effect of task cost and task difficulty on gait.

3.2 Inattentional Blindness and Cognitive Tunnelling

Several studies have examined inattentional blindness and cognitive tunnelling. An early study [43] involved participants that were asked to watch a clip of basketball players passing a ball, with one team in a white uniform and one team in a black uniform. Participants were asked to watch either the white team or the black team and count the number of passes they made. In the middle of the clip, either a gorilla or a woman with an open umbrella walked across the screen. At the end of the study, the participants were asked if they noticed anything odd, and then they were asked directly if they saw the gorilla or the woman if they didn't volunteer the information in the general question. However, in this study the videos used superimposed images of the basketball players. Later, the study was replicated with a more standard modern video clip without superimposing [58]. However, only 44% of participants reported noticing the unusual occurence.

Some studies have looked specifically at inattentional blindness while using a mobile device. In one such study [21], a clown was unicycling in front of a landmark in a busy campus area. Students exiting the area were stopped and asked if they had seen anything unusual, and if they said they hadn't they were asked directly if they saw the clown. Only 25% of cellphone users reported seeing the clown, as opposed to 51% of single people not using cellphones, 61% of people listening to music, and 71% of people walking in pairs. Additionally, even when the clown was not present, observed cellphone users had many more instances of changing directions, performing weaving movements, or stopping, as well as significantly higher collision or near-collision incidents. In a following study [22], cellphone users took longer to avoid a signpost on a path and were less likely to report that they had avoided it or remember what was written on the sign. They were also less likely to notice money hanging from a tree next to the path they were walking on in a second experiment.

In one study, participants drove either while undistracted or while talking on the phone and were afterward asked if they recalled seeing certain objects in the scene. Half as much visual information was reported by distracted drivers than undistracted drivers, even though eye-tracking software was used to confirm the driver had looked directly at the object in question. They also found that talking hands-free while driving was just as distracted as when the phone was held [62].

3.3 Spatiotemporal Gait Literature Review

A literature review was conducted to examine how distraction affects spatiotemporal gait. The results are collected and shown in Table 3.1, and reveals key details on the effect of distractions on spatiotmporal gait. Changes in the velocity of distracted pedestrians seem to be in the range of 4%-42% of baseline. For step length, the range of changes seem to be within 3% to 25% of baseline. For step width, the changes were within 1%-12% of baseline. Changes in step time ranged between 1%-19% of the baseline values. Cadence changed between 3% and 27%. Changes in the speed of people performing different activities on a cellphone while walking were particularly important for the development of the distraction model. Changes in double support time, the time that both feet are on the ground while walking, was less reported but ranged between 17%-112% increases in time, and took up 7%-23% more of the gait cycle on average. Finally, other miscellaneous information was collected, such as how much time the average pedestrian spends on their phone, which was used as a reference for determining distraction intervals in the proposed model.

In the context of crowd simulation, the velocity statistics are important across a wide variety of techniques, including the proposed model. Step width, step length, step time, and cadence may be important for some crowd simulation algorithms and may be beneficial for future work, as well as provide an intuitive understanding of how distraction changes gait.

Table 3.1: Summary of observations found in distraction literature.

Speed (m/s)					
Paper	Walking (Baseline)	Phone Use	Phone Activity	Sample Size	Percent Change
[1]	1.30 $\pm$ 0.12	1.17 $\pm$ 0.10	Text a message describing previous day's activities	18 young adults 20-30 years 8 males/10 females	10.0 $\pm$ 3.8% decrease
[73]	1.70	1.63	Unknown	148 pedestrians 50 distracted pedestrians 98 undistracted pedestrians	4.12% decrease
[57]	1.33 $\pm$ 0.15	Read: 1.16 $\pm$ 0.14 Text:1.01 $\pm$ 0.17	Read: walking at a comfortable pace while reading a passage on a mobile phone screen with minimal manual input other than scrolling through text Text: walking at a comfortable pace while typing the passage 'the quick brown fox jumps	26 adults 29 $\pm$ 11 years 7 males/19 females	Read: 12.78% decrease Text: 24% decrease

			over the lazy dog’		
[6]	Bright light: 1.19 ± 0.12 Dark: 1.18 ± 0.10	Bright/Texting: 1.00 ± 0.12 Dark/Texting: 0.88 ± 0.13 Bright/Gaming: 1.05 ± 0.09 Dark/Gaming: 0.94 ± 0.10	Walking with an obstacle under bright and dark lighting conditions while texting the Korean national anthem or gaming	12 adults 23.50 + 4.89 years 6 males/6 females	Bright/Texting: 15.97 % decrease Dark/Texting: 25.4 % decrease Bright/Gaming: 11.76 % decrease Dark/Gaming: 20.34 % decrease
[24]	1.12 ± 0.23	Downloading: 0.74 ± 0.21 Talking: 0.65 ± 0.26	1) Finding and downloading an Application (App) using a smartphone Or 2) Talking on the phone	26 adults 18-30 years	Downloading: 33.49% decrease Talking: 41.69% decrease
[48]	1.13 ± 0.03	1.04 ± 0.04	Text answers to math problems using words, not numbers E.g. six rather than 6	12 adults 7 males/5 females 28-53 years	7.96% decrease
[50]	0.90 ± 0.10	0.78 ± 0.12	Texting - subjects read and replied to a	18 young adults 25 ± 3 years 7 males/11 females	13.33% decrease

			standard email.		
[70]	1.13 ± 0.17	Texting: 0.95 ±0.18 Texting and listening to music: 0.91 ± 0.17	Texting or texting while listening to music	35 young adults 17 males/18 females 22.09 + 1.63 yrs	Texting: 15.93% decrease Texting while listening to music: 19.47% decrease
[38]	0.78 ± 0.1	Texting: 0.61 ± 0.1 Answering math problems: 0.63 ± 0.1	Texting or answering cognitive math problems without a phone	30 adults 18-50 yrs	Texting: 21.8% decrease Answering math problems: 19.23%
[35]			Texting or talking on the phone	33 healthy participants 13 males/20 females 26 ± 4 yrs	Texting: 33% decrease Talking on the phone:16% decrease
[51]	1.30 ± 0.12	Texting no priority: 1.06 ± 0.18 Texting gait-priority: 1.19 ± 0.15 Texting texting-priority: 0.94 ± 0.17	Texting while walking, and either telling subjects to prioritize texting or walking or not specifying any priority.	32 university students 18-30 yrs	Texting no priority: 18.46% decrease Texting gait-priority: 8.46% decrease Texting texting-priority: 27.69% decrease
[60]			Playing a mobile game	40 young adults 20 males/20 females 19-32 years	27% decrease

				mean height = 176.1±7.5 cm mean weight=69.9 ±10.6 kg	
[53]	Young, Busy hallway: 1.30 ± 0.23 Old, Busy hallway: 1.28 ± 0.12	Walk-look, young, busy hallway: 1.19 ± 0.19 Walk-look, old, busy hallway: 1.20 ± 0.13 Walk-answer, young, busy hallway: 1.11 ± 0.24 Walk-answer, old, busy hallway: 1.01 ± 0.14 Walk-text, young, busy hallway: 1.03 ± 0.19 Walk-text, old, busy hallway: 0.75 ± 0.14	Walk-look: walk and look at a phone, but no other activity Walk-answer: walk and answer questions with no phone Walk-Text: walk and answer questions by text	Young adults between 18-30 yrs and older adults 65 yrs or more	Walk-look, young, busy hallway: 8.46% decrease Walk-look, old, busy hallway: 6.25% decrease Walk-answer, young, busy hallway: 14.62% decrease Walk-answer, old, busy hallway: 21.09% decrease Walk-text, young, busy hallway: 20.77% decrease Walk-text, old, busy hallway: 41.41% decrease
[16]			Walked while texting,	24 undergrad students	Text: 14.86% decrease

			playing a game, or watching a video	12 males/12 females 20.5 yrs avg. age	Video: 12.1% decrease Game: 18.38% decrease
[43]	1.48	Reading: ~1.3 Dialling: ~1.39 Surfing: ~1.23 Selfie: ~1.3	Reading, dialling, surfing, or taking a selfie while walking	36 students 18-30 yrs	Reading: ~12% decrease Dialling: ~6.08% decrease Surfing: ~16.9% decrease Selfie: ~12% decrease
Stride length (m)					
[1]	1.42 ± 0.14	1.34 ± 0.11	Text a message describing previous day's activities	18 young adults 20-30 years 8 males/10 females	5.6 ± 3.5 % decrease
[57]	1.35 ± 0.12	Read: 1.23 ± 0.09 Text: 1.15 ± 0.09	Read: walking at a comfortable pace while reading a passage on a mobile phone screen with minimal manual input other than scrolling through text	26 adults 29 ± 11 years 7 males/19 females	Read: 8.89% decrease Text: 14.8% decrease

			Text: walking at a comfortable pace while typing the passage ‘the quick brown fox jumps over the lazy dog’		
[6]	Bright: 1.36 ± 0.14 Dark: 1.32 ± 0.16	Bright/Texting: 1.22 ± 0.12 Dark/Texting: 1.14 ± 0.14 Bright/Gaming: 1.28 ± 0.11 Dark/Gaming: 1.21 ± 0.12	Walking with an obstacle under bright and dark lighting conditions while texting the Korean national anthem or gaming	12 adults 6 males/6 females 23.50 + 4.89 years	Bright/Texting: 10.14% decrease Dark/Texting: 14.05% decrease Bright/Gaming: 5.88% decrease Dark/Gaming: 8.33% decrease
[24]	1.21 ± 0.16	Downloading: 0.93 ± 0.14 Talking: 0.91 ± 0.18	1) Finding and downloading an Application (App) using a smartphone Or 2) Talking on the phone	26 adults 18-30 years	Downloading: 23.14% decrease Talking: 24.79% decrease
[48]	1.26 ± 0.05	1.20 ± 0.04	Text answers to math problems using words, not numbers	12 adults 28-53 years 7 males/5 females	4.76% decrease

			E.g. six rather than 6		
[70]	1.18 ± 0.13	Texting:1.07± 0.13 Texting and listening to music: 1.05 ± 0.12	Texting or texting while listening to music	35 young adults 17 males/18 females 22.09 + 1.63 yrs	Texting: 9.32% decrease Texting while listening to music: 11.02% decrease
Step length (m)					
[73]	0.86 + 0.10	0.80 + 0.09	Unknown	148 pedestrians 50 distracted pedestrians 98 undistracted pedestrians	6.98% decrease
[6]	Bright: 0.72 + 0.08 Dark: 0.70 + 0.08	Bright/Texting: 0.63 + 0.07 Dark/Texting: 0.59 + 0.06 Bright/Gaming: 0.67 + 0.06 Dark/Gaming: 0.63 + 0.06	Walking with an obstacle under bright and dark lighting conditions while texting the Korean national anthem or gaming	12 adults 23.50 + 4.89 years 6 males/6 females	Bright/Texting: 12.50% decrease Dark/Texting: 15.71% decrease Bright/Gaming: 6.94% decrease Dark/Gaming: 4.9% decrease
[24]	0.60 + 0.08	Downloading: 0.48 + 0.71 Talking: 0.45 + 0.09	1) Finding and downloading an Application (App) using	26 adults 18-30 years	Downloading: 19.95% decrease Talking: 24.6% decrease

			a smartphone Or 2) Talking on the phone		
[50]	0.53 ± 0.06	0.51 ± 0.07	Texting - subjects read and replied to a standard email.	18 young adults 25 ± 3 years 7 males/11 females	3.77%
[38]	0.42 ± 0.3	Texting: 0.36 ± 0.2 Answering math problems: 0.37 ± 0.2	Texting or answering cognitive math problems without a phone	Thirty participants 18-50 yrs	Texting:14.3% decrease Answering math problems: 11.9%
Step width (cm)					
[27]	Young Adults 17.3 ± 2.7 Older Adults 19.6 ± 1.0	Young Adults 19.4 ± 2.9 Older Adults 22.2 ± 1.8	Text: subjects were visually presented a series of ten-digit phone numbers one at a time and asked to dial the numbers on a standard flip phone	16 adults consisting of: 7 younger adults 2 males/5 females 20.4 + 2.2 years 9 older adults 2 males/7 females 61.1 + 10.0 years	Young Adults: 10.8% increase Older Adults: 11.7% increase

[70]	13.40 ± 1.57	Texting: 13.25 ± 1.55 Texting and listening to music: 13.10 ± 1.42	Texting or texting while listening to music	35 young adults 17 males/18 females 22.09 + 1.63 yrs	Texting: 1.12% increase Texting while listening to music: 2.24% increase
Step time (sec)					
[6]	Bright: 0.58 ± 0.04 Dark: 0.58 ± 0.04	Bright/Texting: 0.61 ± 0.05 Dark/Texting: 0.68 ± 0.13 Bright/Gaming: 0.64 ± 0.08 Dark/Gaming: 0.69 ± 0.09	Walking with an obstacle under bright and dark lighting conditions while texting the Korean national anthem or gaming	12 adults 6 males/6 females 23.50 ± 4.89 years	Bright/Texting: 5.17% increase Bright/Texting: 17.24% increase Bright/Gaming: 10.34% increase Bright/Gaming: 18.97% increase
[27]	Young Adults: 0.579 ± 0.049 Older Adults: 0.566 ± 0.037	Young Adults: 0.567 ± 0.032 Older Adults: 0.558 ± 0.054	Text: subjects were visually presented a series of ten-digit phone numbers one at a time and asked to dial the numbers on a standard flip phone	16 adults consisting of: 7 younger adults 2 males/5 females 20.4 + 2.2 years 9 older adults 2 males/7 females 61.1 + 10.0 years	Young Adults: 2.07% increase Older Adults: 1.44% increase

[70]	0.53 ± 0.06	Texting: 0.58 ± 0.08 Texting and listening to music: 0.59 ± 0.09	Texting or texting while listening to music	35 young adults 17 males/18 females 22.09 + 1.63 yrs	Texting: 9.43% increase Texting while listening to music: 11.32% increase
Cadence (steps or strides/min)					
[73] (strides / min)	54.9 ± 2.9	52.4 ± 3.9	Unknown	148 pedestrians 50 distracted pedestrians 98 undistracted pedestrians	4.56% decrease
[6] (steps / min)	Bright: 98.39 ± 6.71 Dark: 97.78 ± 7.07	Bright/Texting: 92.48 ± 7.35 Dark/Texting: 90.55 ± 13.05 Bright/Gaming: 92.80 ± 10.42 Dark/Gaming: 86.18 ± 8.43	Walking with an obstacle under bright and dark lighting conditions while texting the Korean national anthem or gaming	12 adults 23.50 + 4.89 years 6 males/6 females	Bright/Texting: 6.01% decrease Dark/Texting: 7.39% decrease Bright/Gaming: 5.68% decrease Dark/Gaming: 11.86% decrease
[24] (steps / min)	110.32 ± 11.62	Downloading: 90.85 ± 15.54 Talking: 80.65 ± 0.11	1) Finding and downloading an Application (App) using a smartphone	26 adults 18-30 years	Downloading: 17.67% decrease Talking: 26.89%

			Or 2) Talking on the phone		
[48] (steps / min)	108.1 ± 1.5	104.3 ± 1.8	Text answers to math problems using words, not numbers E.g. six rather than 6	12 adults 28-53 years 7 males/5 females	3.52% decrease
[50] (steps / min)	107 ± 10	99 ± 16	Texting - subjects read and replied to a standard email.	18 young adults 7 males/11 females 25 ± 3 years	7.48% decrease
[70] (steps / min)	115.76 ± 9.21	Texting: 103.41 ± 22.00 Texting and listening to music: 103.61 ± 12.73	Texting or texting while listening to music	35 young adults 17 males/18 females 22.09 + 1.63 yrs	Texting:10.67 % decrease Texting while listening to music: 10.50% decrease
[38] (steps / min)	110.4 ± 12	Texting: 100.2 ± 12 Answering math problems: 102.6 ± 12	Texting or answering cognitive math problems without a phone	30 participants 18-50 yrs	Texting: 9.24% decrease Answering math problems: 7.8%
Double Support Time					
[73] (% of the gait cycle)	11.2 ± 2.7	13.3 ± 2.3	Unknown	148 pedestrians 50 distracted pedestrians 98	23% ± 20% increase

				undistracted pedestrians	
[24] (sec)	0.27 ± 0.05	Downloading: 0.49 ± 0.2 Talking: 0.57 ± 0.31	1) Finding and downloading an Application (App) using a smartphone Or 2) Talking on the phone	26 adults 18-30 years	Downloading: 81.48% increase Talking: 111.11%
[70] (% of the walking period)	21.81 ± 2.80	Texting: 23.50±3.71 Texting and listening to music: 23.93±3.57	Texting or texting while listening to music	35 young adults 17 males/18 females 22.09 + 1.63 yrs	Texting: 7.75% increase Texting while listening to music: 9.72% decrease
[38] (sec)	170 ± 30	Texting: 200 ± 30 Answering math problems: 200 ± 40	Texting or answering cognitive math problems without a phone	30 participants 18-50 yrs	Texting: 17.8% decrease Answering math problems: 17.8% decrease
Abs Path Lateral Deviation					
[57] (m)	0.07 ± 0.2	Read: 0.08 ± 0.02 Text: 0.10 ± 0.03	Read: walking at a comfortable pace while reading a passage on a mobile phone screen with minimal	26 adults 29 ± 11 years 7 males/19 females	Read: 14.29% increase Text: 42.86% increase

			manual input other than scrolling through text Text: walking at a comfortable pace while typing the passage 'the quick brown fox jumps over the lazy dog'		
[24] (cm)	1.72 ± 0.00	Downloading: 3.77 ± 2.38 Talking and downloading: 3.83 ± 3.58	1) Finding and downloading an Application (App) using a smartphone Or 2) Talking on the phone	26 adults 18-30 years	Downloading: 119.18% increase Talking and downloading: 122.67% increase
[35] degrees / linear distance travelled	6 degrees 7.5m Linear distance travelled	Texting: 10 degrees Linear distance travelled: 8.5m	Texting or talking on the phone	33 healthy participants 13 males/20 females 26 ± 4 yrs	Texting: 61% increase Talking on the phone: No change
% time spent looking at cell phone while walking					
[70]	0%	70%	Web browsing or texting	32 adults 20-30 years 27 males/5 females	
Other distraction frequency statistics					

[3]	Mean number of texts sent per day: 112 ± 138 = 4.67 ± 5.75 times an hour	128 undergraduate students	
[8]	Self-reported breakdown of how students spend their time: Sleeping: 7.59h Texting: 4.93h On the phone: 4.41h Social Networks: 2.84h Studying: 2.74h Watching TV: 1.94h	534 students 60 males/474 females mean age: 20.8 yrs	
[51]	Texts sent per day on weekdays and weekends: Texts sent % sample on weekday % sample on weekend 0-9 0% 0% 10-30 45.2% 22.6% 31-50 12.9% 32.3% 51-70 25.8% 29.0% 71-90 3.2% 6.5% 91-110 6.5% 3.2% > 110 6.5% 6.5%	32 university students 18-30 yrs	
[51]	Self-reported survey of cell phone use (talking/texting) while walking: on weekday on weekend None 0% 0% < 25% 32.3% 22.6% 25-50% 29% 51.6% 50-75% 29% 25.8% > 75% 9.7% 0%	32 university students 18-30 yrs	
[43]	Self-reported frequency of various distracted activities while walking from 36 participants: Phone texting/reading - often: 14, sometimes: 21, never: 1 Talking on phone - often: 16, sometimes: 19, never: 1 Internet research/browsing - often: 16, sometimes: 8, never: 12	36 university Students 18-30 yrs	

Observed percent of people distracted (not crossing the street)					
[25]		43.5% texting or holding their cell phone 6.8% talking on cell phone 14.1% listening with ear buds 35.6% had no cell phone at all.	Texting or holding phone Talking on phone Listening to ear buds	191 students walking on campus.	
Observed percent of people distracted (at a crosswalk/crossing the street)					
[63]		7.3% text messaging 6.2% using a handheld phone 11.2% listening to music	Texting, using phone, listening to music	Street traffic in Seattle, Washington	
[54]		18.7% distracted	66.5% listening to earbuds ~18% texting or looking down at phone ~8% talking on phone ~7.5% doing more than one of these	High school students crossing the street near their schools	

Chapter 4

Proposed Distraction Model

Some literature on dual-task cost and difficulty was discussed in Section 3.1. We can attempt to classify and categorize activities based on this literature in regards to task costs. An example is shown in Figure 4.1. Talking on the phone, for example, would be a cognitive task exclusively, while reading an article or reading texts from friends would be a primarily visual and cognitive task, and texting is a combination of cognitive, motor, and visual. This figure is meant as an example only, and tasks may be categorized differently depending on the exact nature of the task.

4.1 Fuzzy Goal Concept

Historically, crowd simulations have employed the Principle of Least Effort, first proposed by George K. Zipf [75], to guide agent movement. This is because it produces desirable emergent crowd behaviour [15]. The Principle of Least Effort stipulates that an agent should take the shortest path, moving in a straight line towards their goal and maintains a preferred velocity when possible. However, for distracted pedestrians, the Principle of Least Effort no longer holds as a distracted pedestrian does not consider optimizing their path for minimal effort their top priority. As long as they are mostly making progress towards their target while distracted, they are satisfied.

Studies have shown that walking while distracted produces different gait characteristics, which would not be the case if the Principle of Least Effort always held. For example, studies have shown that some distracted walkers exhibit lateral deviation from a straight path [57, 25, 36]. This is because while the pedestrian is looking down at their device, they may not have an accurate idea

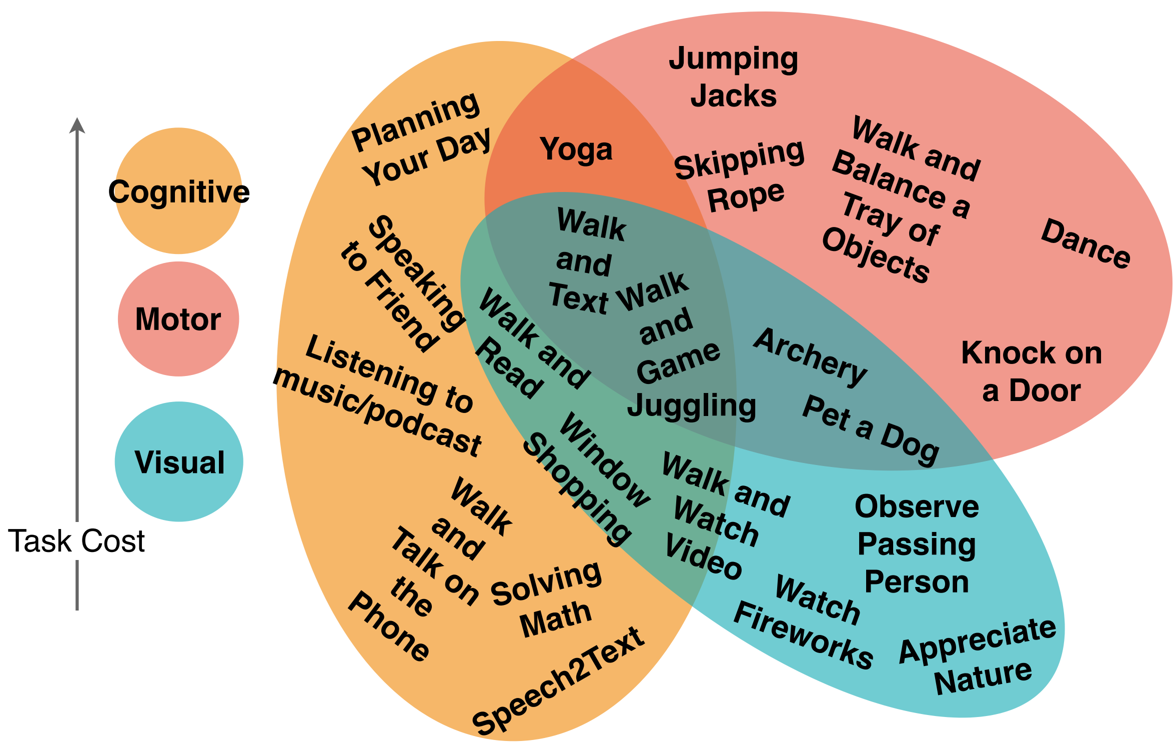

Figure 4.1: An illustration of behaviours categorized into visual, motor, and cognitive loads. A task may demand various types of attention loads. The most distracting tasks tend to be cognitive tasks, followed by motor and then visual. The tasks listed in this figure are for example purposes only, in order to illustrate a way of thinking about task costs in every day life. In reality, tasks such as dancing may be categorized differently depending on the type of dance, etc. Thus there may be subtle nuances of tasks that influence how they are categorized and by whom.

of where their target is ahead of them. They rely on their spatial memory, peripheral vision, and proprioception to guide them towards where they think their goal is. However, these are subpar to paying full attention, which results in clumsy walking patterns. In some cases where the pedestrian is deeply engaged in the distracting activity, they may even forget to turn where they should or continue walking past their target.

This behavior has been modelled with a novel ”fuzzy” goal for distracted agents. Upon becoming distracted, a new fuzzy goal is randomly selected within some ellipsoid of the actual goal, and the agent moves towards this fuzzy goal until they stop being distracted. The size of the ellipsoid is moderated by the distance from the agent to the actual goal or waypoint, which ensures that the fuzzy waypoint is never behind the agent and also means smaller deviations for shorter distraction times which seems desirable. The user can also specify a maximum length of the semimajor and semiminor axes of the ellipsoid. When a fuzzy waypoint is found, a raycast is performed between this point and the actual waypoint to ensure it is not in an unreachable area such as behind a wall. Lateral deviations are modelled if the fuzzy waypoint is laterally deviated from the actual goal or forget-to-turn behaviour is modelled if it is behind the actual goal or subgoal.

In the first set of experiments described in this document, a circular fuzzy goal area is used so the maximum semi-major axis and maximum semi-minor axis are both set to 10% of the agent’s distance to the actual goal. In the second set of experiments, the user can specify these maximum parameters and the simulation will adjust the actual value based on the agent’s distance to the goal.

4.2 Distraction Types and Steering Layer Parameters

Many crowd simulation algorithms use agent parameters that affect agent steering and obstacle avoidance such as their preferred speed, field of view, and neighbour distance. Neighbour distance is the minimum distance between an agent and other agents or obstacles at which they will begin an avoidance maneuver. Based the literature discussed in Section 3.3, these parameters have been adjusted for distracted agents.

Preferred speed of distracted pedestrians has been shown to be decreased in comparison to non-distracted walkers. In addition, the speed decreases more during certain activities than others. For example, texting tends to reduce speed by an average of 20% across various studies [1, 57, 6,

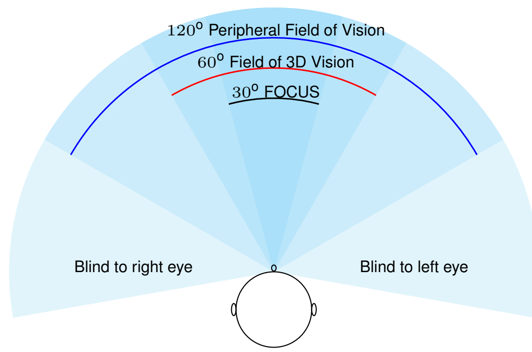

Figure 4.2: An illustration of the key areas of the normative human visual field with their approximate ranges in degrees.

51, 71, 39, 52, 54, 17]. Reading and walking causes speed to be reduced by about 12% [44, 57], and talking on the phone causes about an 18% speed decrease [54].

Another important factor in the path planning of agents is their field of view. A representation of various aspects of the normal human field of view for a nondistracted person is shown in Figure 4.2. The literature suggests that for a distracted walker with gaze directed at their cellphone, not only is the horizontal field of view reduced but the amount of information processed from within that field is reduced. A field of view study done for people talking on the phone suggests that peripheral vision is reduced by 7% for large and 10% for small targets [40] even though the person is still facing forward. Another study showed that the lateral range of eye movement for distracted walkers is reduced to about $18°$ [70]. Therefore field of view for distracted agents is reduced in this model based on the secondary activity they are performing.

While literature regarding neighbour detection distance while distracted is sparse, it is assumed to be reduced for distracted walkers as their visual field is skewed downward and only immediate neighbours are visible in the peripheral vision. For activities where the gaze is directed at a cellphone, the neighbour distance is reduced to 1 meter. For talking on the phone, the gaze is directed forward so the neighbour distance is not reduced. The fuzzy goal radius changes depending on the current distance between the agent and the actual goal, based on lateral deviations reported in the literature [36, 57]. No fuzzy goal is used for talking on the phone behavior as it is assumed there is no deviation of significance when the pedestrian can see their goal. More research in this area may be needed to find optimal values as prior studies have only examined walking in a straight

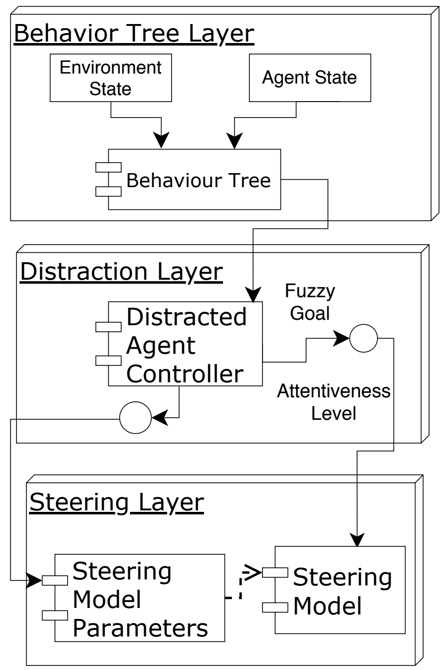

Figure 4.3: A diagram of the module architecture used in the first set of experiments for distracted agents.

line in a very simple environment without interactions.

In the first set of experiments described in the next chapter, the visual attention and refocusing modules are not active. Thus an overview of the distraction model used in these experiments is shown in Figure 4.3. The distraction behaviours are triggered through a behaviour tree. A behaviour tree is a hierarchial tree of tasks often used for AI in robotics, games, and related areas. More complex tasks can be modelled by chaining together groups of smaller tasks, represented as subtrees. All possible actions an agent can take are contained within the behavior tree. A path in the tree from the root to each leaf represents one action, and execution proceeds in a left-to-right manner. If the conditions for one task to execute have not been met, the tree moves on to try the next possible task, and so on [10].

The behaviour tree contains all the actions an agent can perform, arranged as groups of tasks, and controls the logic for triggering when an agent becomes distracted and undistracted. It communicates with the distracted agent controller, which controls distracted behaviour and goal planning including the calculation of the fuzzy goal, which it then sends along with other steering parameters to the steering model.

4.3 A Visual Attention Model For Distracted and Undistracted Agents

A visual attention model for undistracted and distracted agents was developed and acts as a mediator between the agent's extrinsic environment and internal behaviour. It is used to refocus attention based on environmental cues. Attention is a very complicated process that incorporates intrinsic and extrinsic stimuli [66]. To simplify the problem, 3D vision is compressed into a flat 2D vision based on the horizontal field of view, similar to how many previous steering algorithms handle perception of the environment. Our visual attention model takes advantage of the characteristic of human vision where focus is sharpest at the center and becomes increasingly blurrier towards the edges of our peripheral vision. Additionally it has been shown that people focus more on objects and cues closer to the center of their vision and also closer to their person rather than further away [23, 38]. Therefore the proposed model represents a $200°$ sweep of the visual field in polar coodinates. Within this $200°$ sweep two Gaussians are combined to represent the attention attributed to each point, one for the horizontal direction and one for the radial distance. In the horizontal direction, the mean is at the center of the horizontal field and drops off towards the left and right edges. In the radial direction, the mean is at the agent and drops off as distance from the agent increases. These are defined as $\mathcal{N}(0,\sigma_\theta)$ and $\mathcal{N}(0,\sigma_r)$ respectively, where σ_θ is aligned with the forward direction of the agent and $\sigma_r = 0$ similarly is centered on the agent's head. These Gaussian functions are then convolved together to produce an attention distribution across the agent's $200°$ field of view.

This means that we can define a visual attention model for each type of attentive state or behaviour by specifying two parameters, $(\sigma_\theta, \sigma_r)$, and then storing the resulting values in a precomputed texture that can be sampled at runtime. Figure 4.4 shows the visual attention fields represented as textures. Each visual attention model can be stored in a single channel of a texture, allowing for multiple visual attention models to be packed into a single texture for memory efficiency and fast lookup. Values in a texture are normalized. When a lookup occurs, singular vales are sampled using the nearest point on the collider of the agent or obstacle.

The parameters σ_θ and σ_r have been carefully selected for normative vision as well as vision while talking on the phone, reading, and texting, based on reported values in the literature on real-world distracted walkers.

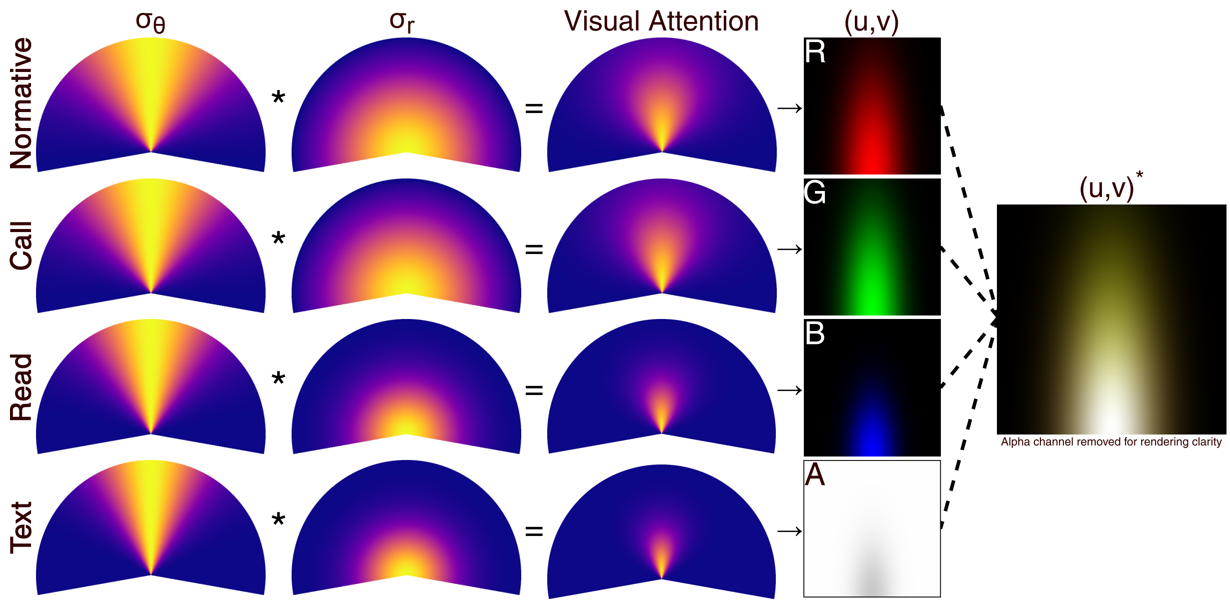

Figure 4.4: An RGBA texture, containing the four visual attention models. Each channel stores a separate visual attention model.

In many steering algorithms, a pedestrian with normal vision is characterized by a neighbour distance [30], planning horizon [59], or query radius [19] with a value between $5-15$m. This range is based on literature on human-human collisions and is often selected with a bias towards lower values [7, 13]. For this reason, the parameter values for a normative agent in this model are set to $\sigma_r = 5m$ and $\sigma_\theta = 30°$. This is consistent with the literature and ensures that the entire range of the visual field is used, although attention is concentrated within $5m$ and within a $60°$ sweep.

Agents talking on the phone are assumed to be able to see as far as an undistracted agent, so again $\sigma_r = 5m$. For σ_θ, it has been shown that visual information is often missed while talking on the phone and also that peripheral vision can shrink artificially due to the cognitive distraction. It has been reported that peripheral vision can shrink $7-10\%$ for large to small targets while talking on the phone. Therefore for this attentive state, σ_θ is reduced by about 7% and becomes $\sigma_\theta = 28°$. The result is that the majority of attention is concentrated within $5m$ and a $56°$ sweep.

The problem becomes more complicated for a reading agent because the head is pitched downward, which significantly affects the vertical range the agent can see in front of them. When people focus their attention on something in their center of vision, it has been shown that their ability to notice and react to information in their peripheral vision decreases [23]. Further, the amount of information loss in the peripheral vision is correlated to the foveal load, the visual complexity of what the person is focusing on, according to the study. However, they also found that subjects who had previously participated in the study had less peripheral information loss than those who hadn't, because they had become "trained" to notice things in their peripheral vision. On average, trained subjects with high foveal load were able to notice 76% of the visual information compared to normative participants with no foveal load. Based on this, σ_θ for readers was set to 76.5% of normative, so for walking and reading agents $\sigma_\theta = 22.95°$. Another study [57] showed that the average head pitch angle for people reading on their phone was $29.22°$ downward. According to a study [9] the worldwide average height at 18 years of age for people born in 1996 was $1.71m$ for men and $1.59m$ for women, so an average of these two values, $1.65m$, was used to calculate a focus distance of $\sigma_r = 2.95m$ on the ground in front of them. So attention is then focused within a $46°$ sweep and $2.95m$.

Similarly, a texting agent also has reduced field of view while looking down at a cellphone. Again, it is assumed that agents have a high foveal load and are trained to notice things in their peripheral vision. However, as previously discussed, texting has an additional motor demand and may have higher cognitive demand than reading. So, an additional 10% peripheral visual information processing loss is estimated. Then σ_θ is 66.5% of normative, so $\sigma_\theta = 19.95°$. However, it has been reported that head pitch angle while texting tends to be slightly greater than while reading [57] at about $31.80°$, which can be used with an average human height value of $1.65m$ to calculate $\sigma_r = 2.66m$ in a similar fashion as for reading. Therefore, attention is focused within $2.66m$ and a $40°$ sweep.

4.4 Refocusing

A probabilistic refocusing model was developed to determine when an agent should stop being distracted and look up. It produces one of two possible outcomes in *refocus*, *continue*. When an agent is currently distracted, at a user specified frequency f_r the reaction model executes three possible Bernoulli trials. In these experiments, $f_r = 3$Hz which is approximately the average human reaction time. However, more or less responsive pedestrians can be produced by adjusting this value lower or higher respectively.

The first of the three Bernoulli trials executed by the reaction model represents intrinsic properties of the agent and determines if the agent should look up based on the agent's personality or situation. The value of $Pr(\textit{refocus}) = p_i$ may represent how likely a person will prioritize paying full attention to where they're going over performing cellphone activities such as taking important calls or responding to text messages from friends, which is dependent upon their personality. This value is user-specified and can be modified per agent according to the user's needs. The value remains fixed, and this Bernoulli trial executes at every reaction check, so it is part of a Bernoulli process and can be presented as a Binomial distribution $\mathcal{B}(N_r, p_i)$ where N_r is the number of reaction checks $N_r = T_d f_r$ and T_d is the total time the agent has been distracted up to the timestep $t \in T$.

The second and third Bernoulli trials are designed to represent extrinsic aspects of other objects or people in the agents environment. These trials depend on objects (and other agents) and their locations in the agent's visual field. These trials are run for each object or agent in the field of view during every refocus update. The first of these is based on an extrinsic property of the object in the field of view that represents the visual appeal of the object and determines whether the agent will notice the object and refocus their attention. The value of $Pr(\textit{refocus}) = p_e$, is intended to represent an object's color, shape, contrast, or other interesting properties of the object. The probability that an object's visual appeal catches the attention of the agent and causes them to refocus $Pr(\textit{refocus}) = p_e$ is conditioned on the visual attention field as follows: $P(\textit{refocus}|p_e, r, \theta)$. Similarly, the third trial is designed to represent an object's eye-catchingness due to its relative speed, which is known to have a significant impact collision prediction [31], and so is considered separately in these experiments. The value

$$Pr(\textit{refocus}) = p_v = \frac{|v_{rel}|}{2|v_{\max}|} \tag{4.1}$$

determines whether an agent will refocus based on an object's relative movement to their own, where $v_{\max}$ is the maximum velocity of any object in the scene. This is also condition on the visual attention field and the object's location within that field such that $P(\textit{refocus}|p_v, r, \theta)$.

A random task time is selected when a distracted activity is triggered. The minimum and maximum task time may be specified by the practitioner. If a refocus event occurs while distracted, the agent will pause their activity and look up, then after another randomly selected wait time they

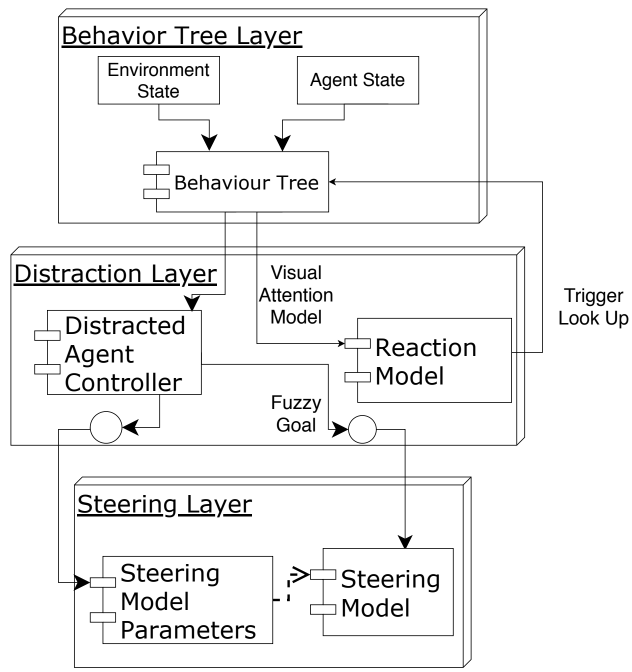

Figure 4.5: A diagram of the module architecture including the reaction model used in the second set of experiments for distracted agents.

will resume the task as long as at least 1 second of task time remains, to prevent oscillations.

If any of the Bernoulli trials succeeds, the agent looks up and the visual attention model switches to normative. For reading and texting behaviours the agent will interrupt their activity and lower their device. However, for the talking behaviour the activity is not interrupted and only the visual attention model becomes normative. This is meant to model pedestrians talking on the phone who do not typically hang up on a call, but may pay more attention to their surroundings if they've noticed something interesting.

The previous model hierarchy was shown previously in Figure 4.3. The reaction model sits in the distraction layer and communicates with the behaviour tree. The behaviour tree sets the current visual attention model in the reaction model as well as the reaction parameters like p_i. The reaction model performs the reaction checks and if any succeed then it sends a trigger look up event back to the behaviour tree. The behaviour tree then triggers the look up and sends the appropriate parameters to the distracted agent controller and steering model. A new representation of the model hierarchy including the reaction model is shown in Figure 4.5 .

4.5 Crowd Authoring

In order to test these methods with virtual crowds quickly and easily, an intuitive crowd authoring tool was developed for Unity3D. The customization interface is shown in Figure 4.6. Each scenario is set up with one or more crowd spawning areas around the scene, and one or more goals. Crowd spawning areas are simply planes that encompass the area within which agents can be spawned. The desired spawning area and goal can be dragged by the user from the scene hierarchy into the floor and goal fields respectively. The user can then specify how many of each type of agent should be specified and many other options for distracted agents. Normative agents will not become distracted, while distracted agents have cellphones and may become distracted and undistracted according to the specified parameters and the visual attention model. Desired distraction parameters can be set by opening up the distraction parameters foldout, where the user can then set probabilities for various distracted activities, minimum and maximum distraction task time, etc. Once the user is satisfied with the parameters, clicking the Create Agents button will spawn the agents into the scene and they will appear in the scene view and hierarchy. The user can then further tweak positions or other parameters as they see fit. Groups of agents can be spawned in multiple spawning areas or in the same spawning area as other groups, and with the same or different goals as other groups. A similar crowd generation tool has also been implemented for spawning crowds at runtime, where agent locations are randomized at every runtime.

When agents are generated, the navmesh is sampled for a free location within the spawn region. Once a location is found, there is a check to make sure there is sufficient distance to any nearby agents, and then the agent is spawned. Obstacles can be spawned in a similar fashion, and the navmesh will update automatically. Parameters such as an obstacle's extrinsic probability, p_e, can be set before obstacles are spawned and allows the user to control the visual appeal or eye-catchingness of an object. The p_e of an agent can also be set before spawning agents which affects agents around them, as well as their own p_i and their f_r.

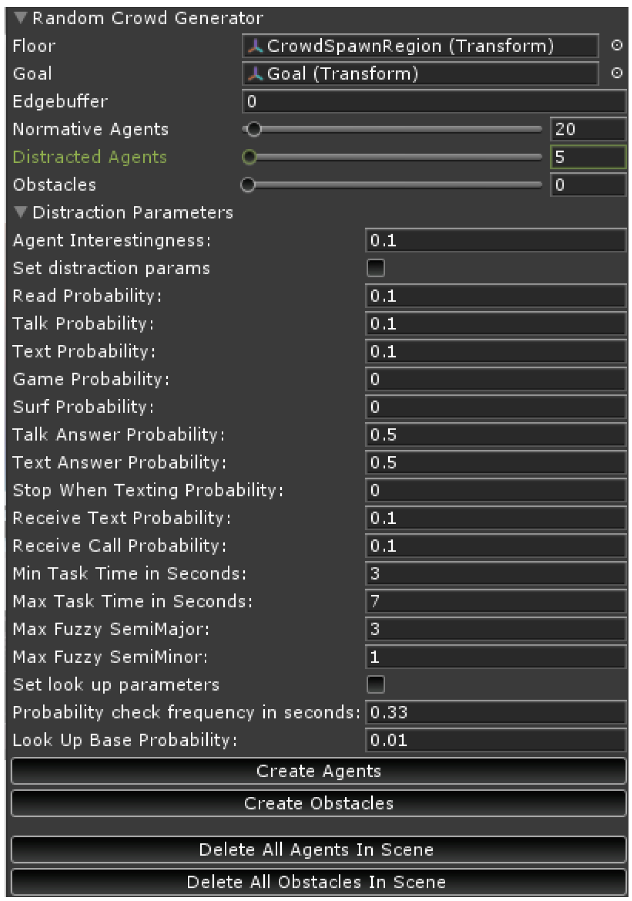

Figure 4.6: The crowd authoring tool interface with a number of customizable parameters.

Chapter 5

Preliminary Experiments Without An Attention Model

A set of preliminary experiments demonstrated the effect of three types of cellphone distractions on crowd flow statistics. The three types of distractions were texting, talking, and reading on a cellphone while walking. In both sets of experiments, PAM was used as the steering algorithm. In this set of preliminary experiments, the built-in vision model of PAM was used. The experimental setup and results for the preliminary experiments is described below.

5.1 Experiments

Pedestrians performing three types of cellphone distractions were studied in simulation, which were texting, talking on the phone, and reading on a cellphone. While distraction types are usually triggered based on events defined in a behavior tree, for the purposes of this experiment the distraction was triggered immediately at the start of the simulation and remained active for the duration of the trial, so there were no state changes. Additionally in any trial, only a single distraction type could be active for all distracted agents in the simulation, so that comparisons could be made between trials examining different distraction types. Finally, scenarios included Egress and Crossing Groups scenarios to study quantitative results as these are scenarios traditionally used in large crowd simulation studies, and also several simple scenarios to demonstrate qualitative effects. In these experiments, the visual attention and refocusing models were not used. To demonstrate

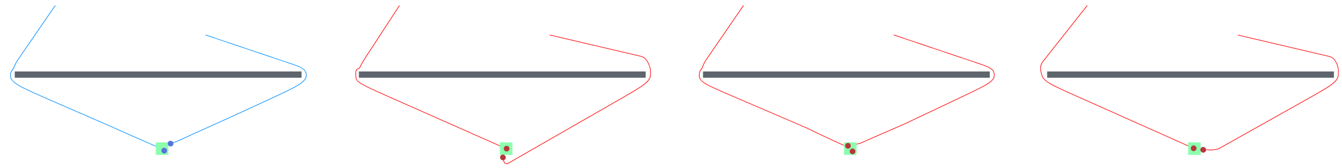

(a) (b) (c) (d)

Figure 5.1: The trajectories in the simple wall scenario of (a) two normative agents, (b) two distracted agents reading, (c) two distracted agents talking on phone, and (d) two distracted agents texting.

intuitive aspects of the model, the qualitative effects will be described first.

5.1.1 Qualitative Impacts

To demonstrate the intuitive aspects of the distraction model, a number of simple scenarios commonly used to test robustness of steering algorithms were tested. These included the simple wall, oncoming obstacle, surprise, and diametric goals scenarios. High quality screen captures with trajectory renderings were taken to compare between undistracted (normative) agents and the three distraction types. Trajectory comparisons for the Simple Wall scenario are shown in Figure 5.1. Trajectory comparisons for the oncoming obstacle environment can be seen in Figure 5.2. Trajectory comparisons for the surprise environment can be seen in Figure 5.3. Trajectory comparisons for the diametric goals environment can be seen in Figure 5.4.

5.1.2 Egress

Chapter 4 discussed how research from the literature was used to build a model of distraction for individual agents. The preliminary scenarios discussed in Section 5.1.1 were designed to show how this model affects steering behavior for individual agents, small interactions between agents, or small groups. However, in crowd simulation it has often been shown that emergent crowd behavior can manifest from microscopic models designed for individual agents. Then the model can be tested in a large crowd scenario and observe what emergent effects occur. In particular, what kind of impact might distraction behaviours have on bottleneck egress openings which are commonly found in public event halls and arenas and are a traditionally tested environment in crowd simulation literature. For this experiment, several crowd simulation statistical measures including flow rate,

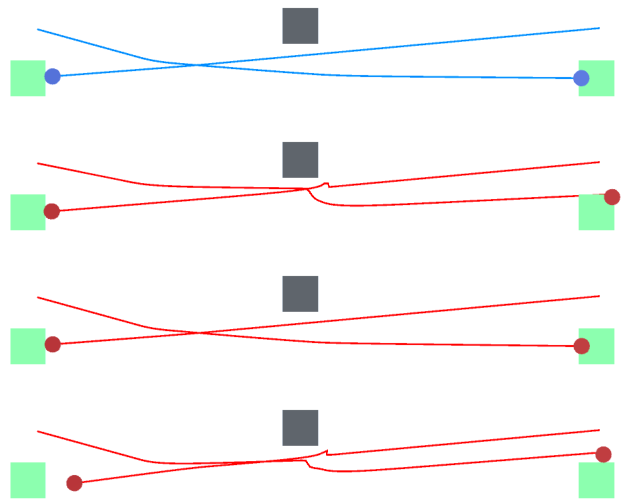

Figure 5.2: The trajectories in the oncoming obstacle scenario of (from top to bottom): two normative agents, two distracted agents reading, two distracted agents talking on the phone, and two distracted agents texting.

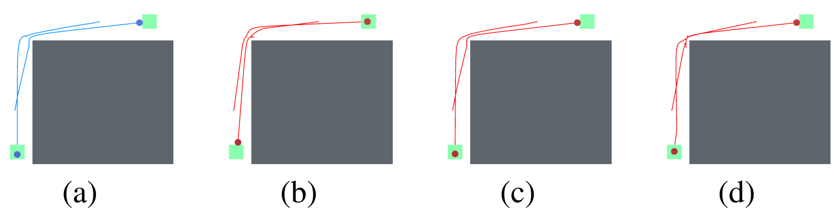

Figure 5.3: The trajectories in the surprise scenario of (a) two normative agents, (b) two distracted agents reading, (c) two distracted agents talking on phone, and (d) two distracted agents texting.

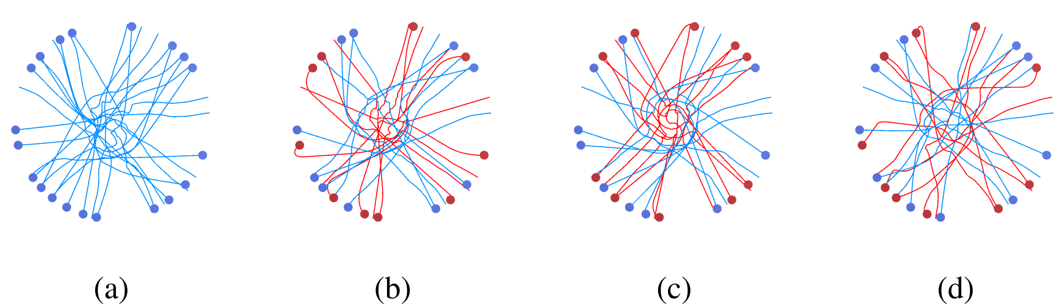

Figure 5.4: The trajectories in the diametric goals scenario of (a) 20 normative agents, (b) 10 normative and 10 distracted agents reading, (c) 10 normative and 10 distracted agents talking on phone, and (d) 10 normative and 10 distracted agents texting.

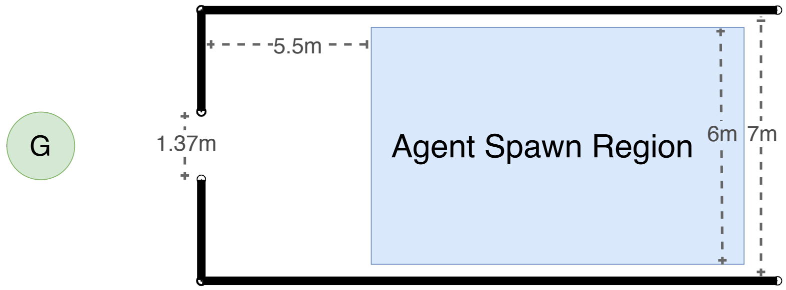

Figure 5.5: An overview of the egress environment. All agents start in the hallway and head towards a shared goal beyond a single bottleneck point.

effort, and kinetic energy will be observed. The hypothesis for this experiment is that distraction behaviours including texting, reading, or talking on the phone have an adverse effect on these crowd simulation metrics. Afterward, quantitative analysis will be done again on high-quality screen captures to assess any qualitative artefacts that may appear when distracted agents are present in the simulation.

The layout of the Egress scenario is shown in Figure 5.5. Each of the texting, talking, and reading distraction types were examined separately such that all distracted agents in any given trial were performing the same distraction type. Every trial consisted of a total of 50 undistracted and distracted agents altogether. Then for each distraction type, successive batches of trials were run such that the ratio of distracted agents performing that activity increased with every batch of trials while the number of undistracted agents decreased. For example, the first set of trials had 50 undistracted agents and 0 distracted agents, the next set had 49 undistracted agents and 1 distracted agent and so on. Trials therefore had 0, 1, 5, 10, 20, 30, 40, or 50 distracted agents and for each of these numbers 20 trials with randomized initial conditions were conducted. Remaining agents in these randomized trials were undistracted agents who used the default parameters of PAM and a preferred speed of 1.3m/s which is typically considered the average walking speed of people. Path length, effort, and kinetic energy were measured during each trial. Path length is the sum of all agent trajectories represented as multi-segmented curves. Effort is an estimated total of metabollic energy spent over the path integrall of all agents. Kinetic energy is defined a the total kinetic energy over the path integral of all agents. Results are analyzed by finding the mean and standard deviation of these collected measures across agent models and ratio proportions of the total crowd and conducting a Kruskal-Wallis sign rank sum test. Post-hoc Conover's test with FDR and Holm corrections were

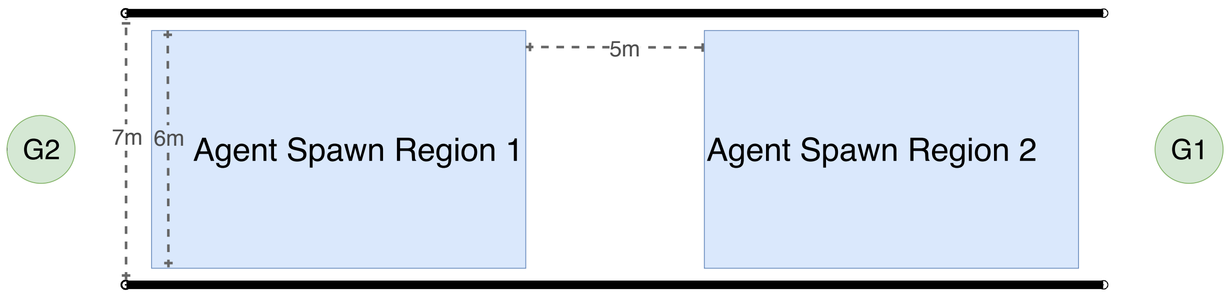

Figure 5.6: An overview of the crossing groups environment. Agents start in groups at opposite ends of a hallway with respective goals at the opposite ends past the opposing group.

also used to identify significant dominances between the models across proportions of the crowd.

5.1.3 Crossing Groups

In Section 5.1.2, it was discussed how interesting individual agent behavior models could potentially have an emergent impact on large-scale agent-environment and agent-agent interactions. To investigate this, distracted agents were observed in the egress bottleneck scenario. A characteristic of the egress scenario is that all agents are heading in the same direction, so there is less agent-agent interference. This is so that the impact of the bottleneck environment can be isolated as much as possible. However, agent-agent interference is also a subject of interest. How much do distracted pedestrians in the real world disrupt the walking systems commonly found in real-world areas? For example, how much do distracted pedestrians interfere with laminar flow and cause congestion, and to what extent do they hinder people who are not distracted? With these questions in mind, we take a look at distracted pedestrians in the Crossing Groups scenario. In this scenario, two groups of agents going in opposite directions must pass through one another while avoiding collisions. In this experiment, the hypothesis is that distracted behaviours such as texting, talking on the phone, or reading has a negative impact on crowd statistical measures such as path length, effort, and total kinetic energy.

The layout of the Crossing Groups scenario is shown in Figure crossing-groups-scenario. Similar to the Egress scenario, each of the distraction types were tested in isolation so that all distracted agents in any given trial used the same distraction model type. Sets of 20 trial batches were performed for each distraction type where the number of distracted agents out of a total of 50 agents increased in each batch, resulting in trials with 0, 1, 5, 10, 20, 30, 40, and 50 distracted agents out of

50 total agents with the remaining agents being undistracted normative agents. Undistracted agents used default parameters of PAM with a desired speed of 1.3 m/s and on every trial initial conditions were randomized. Path length, effort, and kinetic energy were again measured and compared between models and across proportions of distracted agents in the crowd. Mean and standard deviation were derived from the collected data and then analyzed with a Kruskal-Wallis sign rank test, post-hoc Conover's test with FDR and Holm corrections to identify significant dominances.

5.2 Results and Discussion

This section will present the results of the preliminary experiments and a discussion. First, the qualitative screen captures of the simple scenarios will be examined, and then the quantitative data from the Egress and Crossing Groups scenarios will be discussed.

5.2.1 Qualitative Results

The screen captures show clear differences between the distraction types for each scenario. In all scenarios, the path of least effort is taken by the normative agents as they head toward their goal as efficiently as possible. Wider turns are taken in by the reading and texting agents, making their trajectories less efficient. They may end up past or deviated from the goal and have to make an adjustment, which is sometimes observed by real-world pedestrians. Talking agents tend to exhibit trajectories close to normative behavior because they do not exhibit lateral deviation and their trajectory is not effected by their slower speed and slightly reduced field of view when there are no other agents in their vicinity.

In the oncoming obstacle environment, the two normative agents avoided each other without issue. However, texting and reading agents both failed to avoid both each other and the obstacle, indicated by the bump in the trajectories where the obstacle was hit. Again the talking behavior looks similar to normative agents.

In the surprise scenario, the two agents had a collision at the corner when they were reading or texting, and the collision had greater impact for texting agents. Undistracted agents had no trouble avoiding each other, as well as talking agents, probably because field of view is only slightly reduced for talking agents.

Table 5.1: A table showing significant differences in dominance in crowd simulation metrics for pairs of models and the number of distracted agents out of a total of 50 agents in the egress scenario. Significant differences are indicated by checkmarks ($p < 0.05$).

	Flow Rate								**Effort**								**Kinetic Energy**							
	1	2	5	10	20	30	40	50	1	2	5	10	20	30	40	50	1	2	5	10	20	30	40	50
No Distraction/Read	X	X	X	X	X	X	✓	✓	X	X	X	X	X	X	✓	✓	X	X	X	X	X	X	X	X
No Distraction/Talk	X	X	X	✓	✓	✓	✓	✓	X	X	X	X	X	X	✓	✓	X	X	X	✓	✓	✓	✓	✓
No Distraction/Text	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	✓	✓	✓	✓	✓
Read/Talk	X	X	X	✓	✓	✓	✓	✓	X	X	X	X	X	X	✓	✓	X	X	X	✓	✓	✓	✓	✓
Read/Text	X	X	X	X	X	X	✓	✓	X	X	X	X	X	X	✓	✓	X	X	X	X	✓	✓	✓	✓
Talk/Text	X	X	X	✓	✓	✓	✓	✓	X	X	X	X	X	X	✓	✓	X	X	X	✓	✓	✓	✓	✓

In the diametric goals scenario, the trajectories are most entangled at the center for the reading behavior, while paths appear smoother for the texting behavior. The paths of the talking agents seem even smoother than that of the undistracted agents. The reason for these perhaps unintuitive results lies in the speed of the agents. When all agents are going the same speed as for undistracted agents, they all meet in the middle at once and must navigate around each other, whereas when some agents are talking and have reduced speed, by the time the talking agents reach the center some of the normative agents have already made their way out of the center, so there is less congestion at the center and therefore smoother trajectories. The same is true for talking and reading agents. Since talking agents are even slower than reading agents, by the time they reach the center even more of the congestion has cleared.

5.2.2 Egress Results

Results for the Egress scenario are summarized in Figure 5.7 which shows the crowd metrics across all models and proportions. Table 5.1 summarizes the results from the statistical tests, showing at what number of distracted agents the models became significantly different from each other. For example, when looking at flow rate, the reading behaviour became significantly different from no distraction when there were at least 40 reading agents, and so on.

The results of the experiment show interesting effects of distractraction modalities in the bottleneck Egress scenario. While texting does not particularly impact flow rate or expended effort, it has a significant impact on kinetic energy when the proportion of distracted agents in the crowd

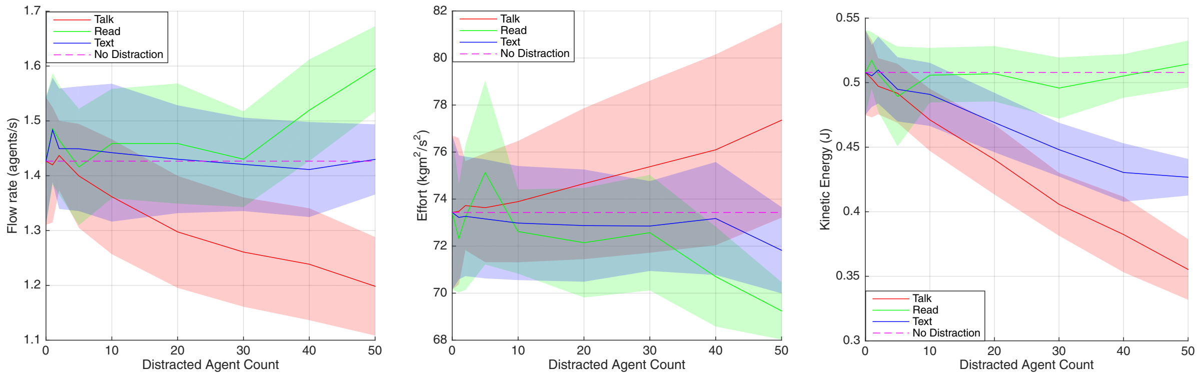

(a) (b) (c)

Figure 5.7: Graphs of the quantitative analysis for the four agent models and the proportion of distracted agents in the egress scenario for (a) flow rate, (b) effort, and (c) kinetic energy.

Table 5.2: A table showing significant differences in dominance in crowd simulation metrics for pairs of models and the number of distracted agents out of a total of 50 agents in the crossing groups scenario. Significant differences are indicated by checkmarks ($p < 0.05$).

	Flow Rate								**Effort**								**Kinetic Energy**							
	1	2	5	10	20	30	40	50	1	2	5	10	20	30	40	50	1	2	5	10	20	30	40	50
No Distraction/Read	X	✓	✓	✓	✓	✓	✓	✓	X	X	✓	✓	✓	✓	✓	✓	X	✓	✓	✓	✓	✓	✓	✓
No Distraction/Talk	✓	✓	✓	✓	✓	✓	✓	✓	X	X	X	✓	✓	✓	✓	✓	X	X	✓	✓	✓	✓	✓	✓
No Distraction/Text	✓	✓	✓	✓	✓	✓	✓	✓	X	X	✓	✓	✓	✓	✓	✓	X	✓	✓	✓	✓	✓	✓	✓
Read/Talk	X	✓	X	✓	✓	✓	✓	✓	X	✓	✓	✓	✓	✓	✓	✓	X	✓	✓	✓	✓	✓	✓	✓
Read/Text	✓	✓	✓	✓	X	✓	X	✓	X	X	X	X	X	X	X	✓	X	X	✓	✓	✓	✓	✓	✓
Talk/Text	X	✓	✓	✓	✓	✓	✓	✓	X	✓	✓	✓	✓	✓	✓	✓	X	✓	✓	✓	✓	✓	✓	✓

reaches 20%. Talking on the phone had a significant impact on flow rate and kinetic energy above 20% distracted agents. Thus it can be concluded that when the proportion of distracted agents in the crowd reaches 20%, flow rate and kinetic energy are adversely impacted in mode dependent ways.

5.2.3 Crossing Groups Results

Figure 5.8 shows the plotted results containing the metrics across all models and their proportions. Table 5.2 summarizes the results of the statistical tests for the Crossing Groups scenario. Qualitative examples in the Crossing Groups environment can be seen in Figure 5.9 where agent paths are shown.

The results generated from the Crossing Groups scenario are especially interesting. The extent

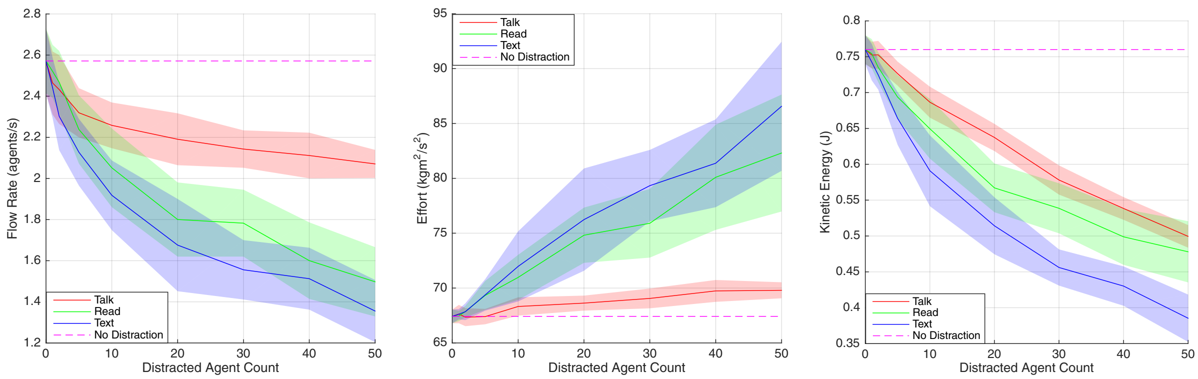

(a) (b) (c)

Figure 5.8: Graphs of the quantitative analysis for the four agent models and the proportion of distracted agents in the crossing groups scenario for (a) flow rate, (b) effort, and (c) kinetic energy.

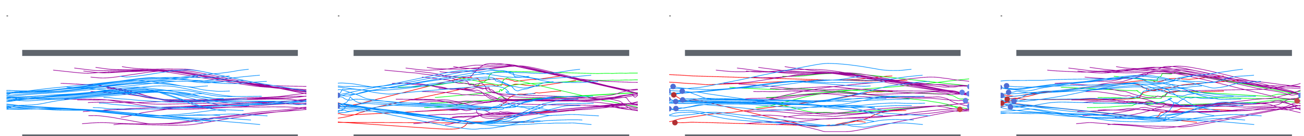

(a) (b) (c) (d)

Figure 5.9: The trajectories in the crossing groups scenario of (a) 50 normative agents, (b) 40 normative and 10 distracted agents reading, (c) 40 normative and 10 distracted agents talking on the phone, and (d) 40 normative and 10 distracted agents texting.

(a) (b)

Figure 5.10: The trajectories in the simple wall scenario of (a) 50 normative agents, (b) 40 normative agents and 10 distracted agents texting. In both cases, the initial positions and conditions were the same. The inclusion of distracted agents disrupts the laminar flow.

to which laminar flow is disrupted by distracted pedestrians far exceeds expectations. An example is shown in Figure 5.10. Even a single distracted agent in a crowd of 50 exhibiting the talking on the phone or texting behaviours can significantly impact the flow rate of crossing groups. The same is true for the reading behavior at 2 agents and all distracted behaviors in this paper at 5 agents out of 50. Across distraction types, each type significantly differs from each other after around 2 out of 50 agents. This can be seen in Figure 5.9 which shows that agent paths are much smoother when no distracted agents are present. Distracted agents disrupt lane formation behaviors commonly observed in the real world which slows down other agents as well. These results show that simulating distracted pedestrians in simulations is important and that different types of distractions can impact crowd flow in remarkably different ways, which has important implications in our world where distracted pedestrians on their cellphones is an increasingly common sight.

Chapter 6

Experiments With An Attention Model

Following the preliminary experiments, the model was extended to include the visual attention and reaction models described in Chapter 4, replacing the vision model of PAM. A second set of experiments was conducted to evaluate the effects of these additional models on crowd authoring. This chapter describes the experimental setup and results for these experiments.

6.1 Experiments

Experiments were conducted to evaluate the effects of the visual attention model, refocus model, and fuzzy goal ellipsoid model and their parameters. Additionally, computational performance was measured. Finally, some qualitative results were collected and examined.

Computational performance was measured first. This experiment was carried out in an open environment in which 100 obstacles are generated and randomly placed in a grid like fashion, and then agents and their goals are randomly placed at each trial. An example configuration of this environment in shown in Figure 6.1. In this environment, the number of distracted agents was gradually increased and the frame time in milliseconds was measured. A GPU-based system using a NVIDIA GeForce GTX 1060 and an Intel i7-7700HQ CPU was used. The experiment was conducted using Unity in Release mode.

To illustrate the effect of the refocus model, each parameter was varied in isolation while the other parameters remained static. First, the intrinsic probability p_i was examined in a scenario space environment containing a fixed number of 20 distracted agents. The p_i parameter was slowly

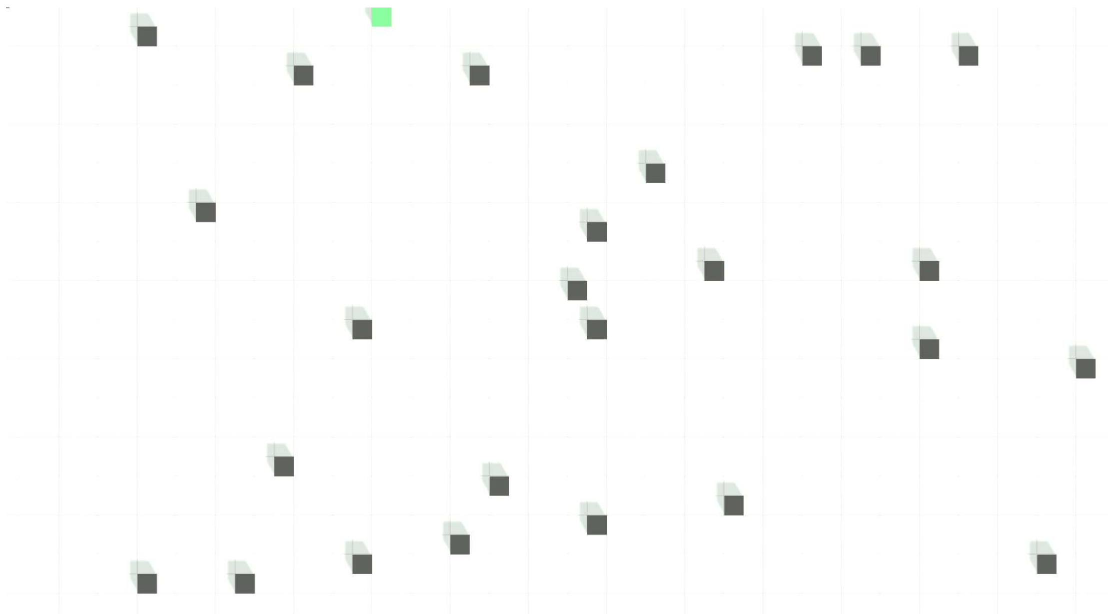

Figure 6.1: A possible configuration of the Scenario Space environment.

increased from 0 to 1 across trials but was homogeneous across all agents for any given trial. For each trial, the number of refocus events was counted.

A similar process is used to demonstrate the effect of the p_e value. Again the scenario space environment was used. For 20 homogeneous distracted agents, the p_e factor for all obstacles is slowly increased from 0 to 1 and refocus events were counted.

Finally, the effect of the p_v parameter was evaluated, this time using a unidirectional and bidirectional flow scenario. In a unidirectional flow scenario, all agents are moving in approximately the same direction, so the relative velocity between agents is going to be low in comparison to a bidirectional flow scenario where groups are moving in opposite directions. In both scenarios, refocus events are counted and compared to evaluate the impact of the p_v.

The impact of the visual attention model for each of the distraction types in demonstrated in the scenario space environment with 20 agents. All agents in each trial are homogeneous, however for each experiment the distraction type and visual attention model is restricted so that only one of the four types may be active - normative, reading, talking, or texting. For each distraction type, we measure refocus events, collisions, and average path length. Then the distraction types are compared against each other based on these measurements.

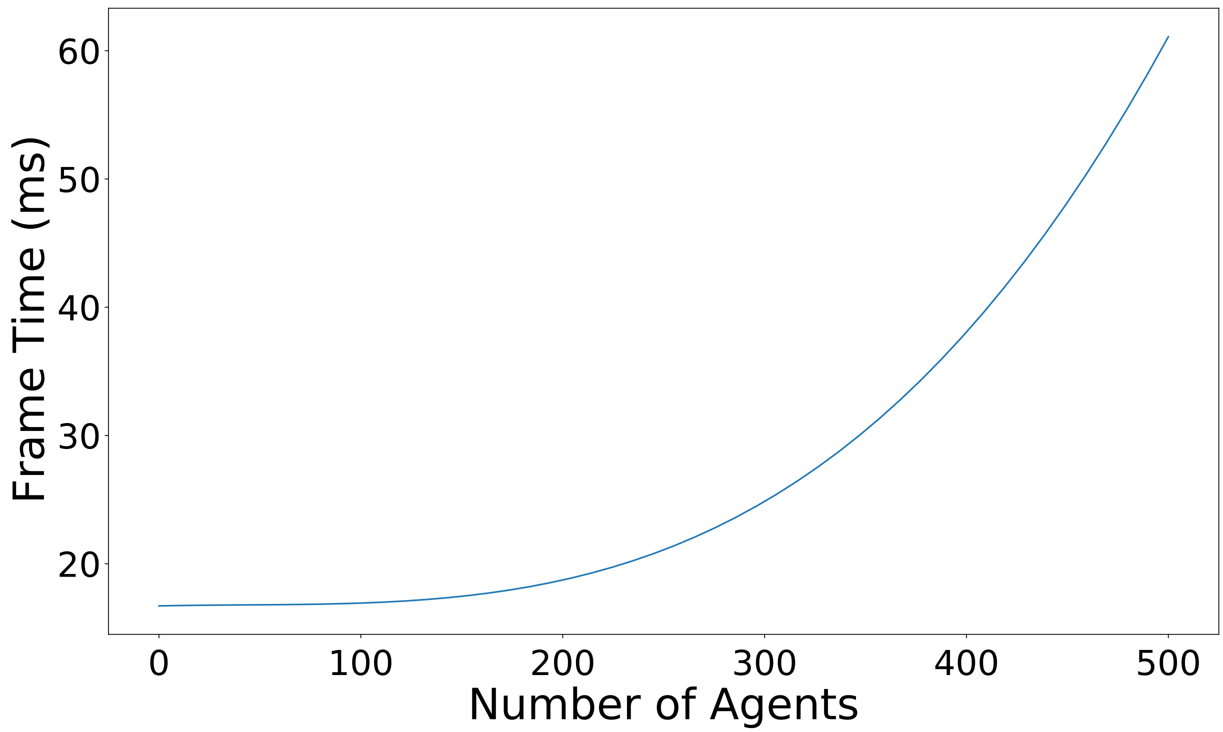

Figure 6.2: A graph showing the increase in frame time in milliseconds as the number of distracted agents increases.

To evaluate the impact of the fuzzy goal parameters, a single agent in the scenario space environment is used. At each trial, random values for the maximum semi-major and semi-minor fuzzy goal ellipse lengths are selected. Here, random values are chosen between 0 and 1, but higher values for these parameters are also possible. A random goal for the agent is also selected. Then the path length is measured for the agent.

6.2 Results and Discussion

6.2.1 Computational Performance

The results for computational performance in the scenario space environment are shown in Figure 6.2, which shows the frame time in milliseconds as the number of distracted agents increases. The system can support more than 200 distracted agents without a significant increase in frame time. However, as the number of distracted agents increases towards 500, the frame time begins to increase. This may improve with optimization of the implementation.

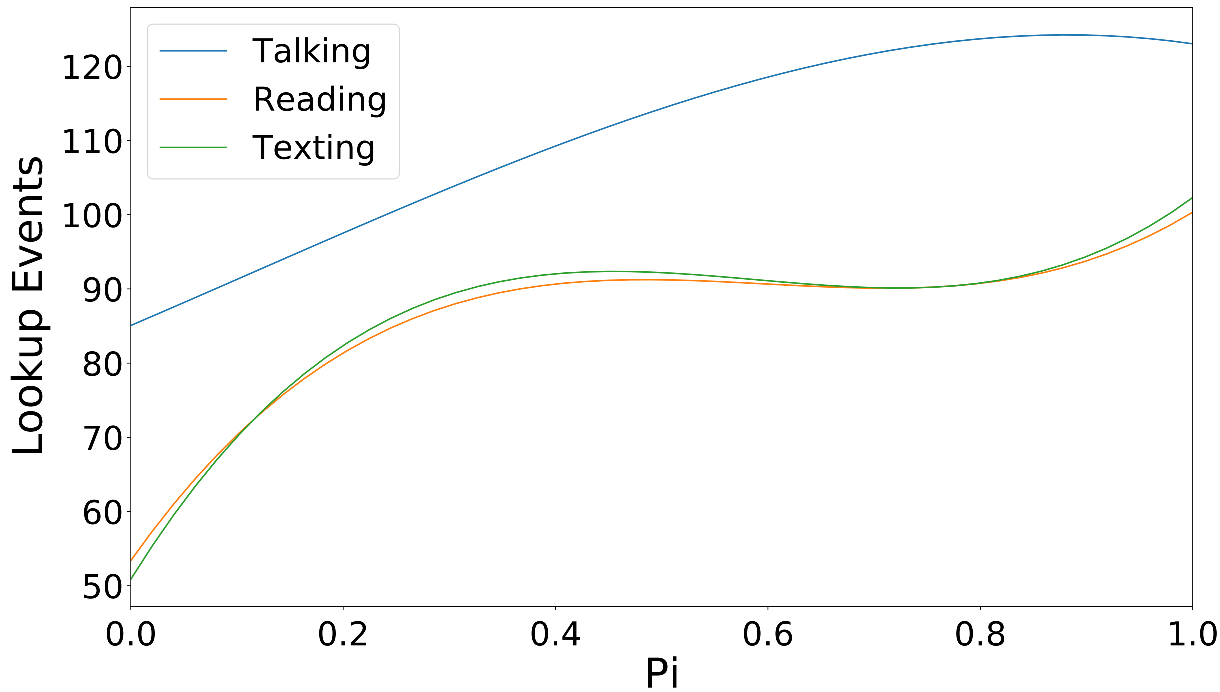

Figure 6.3: The number of refocus events of each distraction model as p_i values increase.

6.2.2 Refocus Model Parameters

Results showing refocus events for increasing values of p_i are shown in Figure 6.3. As expected, higher values of the p_i result in higher numbers of look ups. Talking in particular shows higher numbers of refocus events in general than other distracted behaviours, which is expected, because the visual field for talking agents is larger. There may be more objects or agents in talking agents' field of view and they remain in the field of view for longer, so the chances of a look up being triggered by the p_e and p_v factors is higher.

Similarly, results showing refocus events for increasing values of p_e are shown in Figure 6.4. As expected, once again higher values correpond to higher numbers of refocus events.

As the p_i and p_e values increase, their impact on refocus events plateaus. This is because of the high frequency of the probability checks. Since a frequency of 3 times per second was used, even small increases in the p_i and p_e are enough to trigger a look up behaviour most of the time, so further increases only have a small impact.

Finally, results showing refocus counts in uni-directional and bi-directional flow environments are shown in Figure 6.5. As expected, the number of look up behaviours was slightly higher for

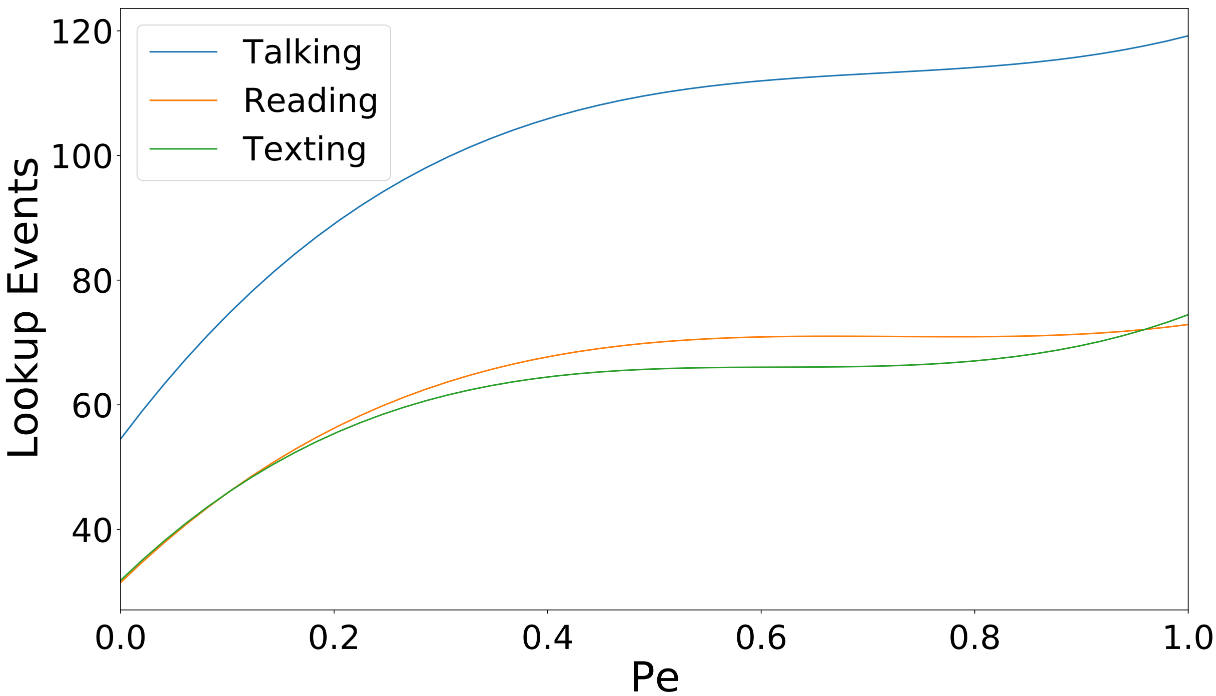

Figure 6.4: The number of refocus events of each distraction model as p_e values increase.

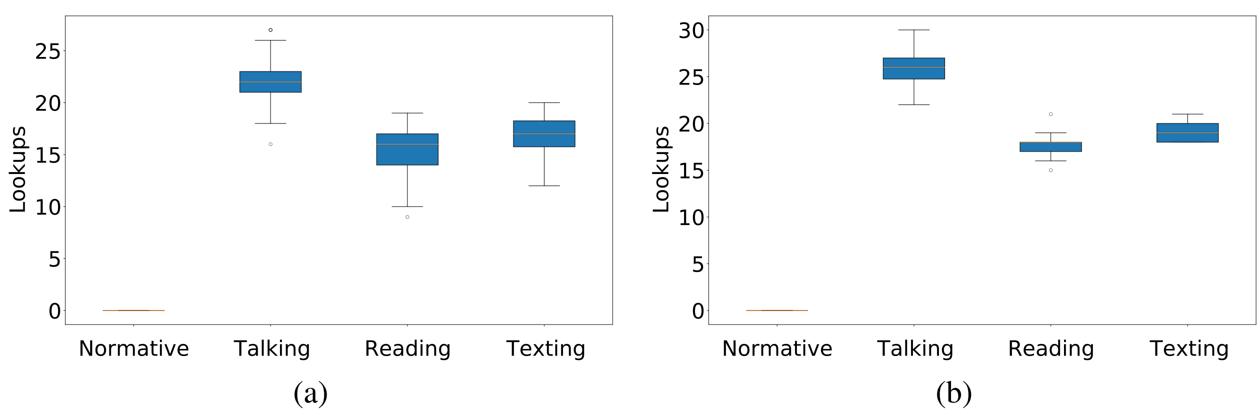

Figure 6.5: The number of focus events for the different distraction models in (a) uni-directional and (b) bi-directional groups.

the bi-directional group environment. Texting agents had slightly more look up behaviours than reading agents, which is somewhat unexpected. A possible explanation may be that because texting agents move slower than reading agents, other objects and agents may remain in their visual fields for longer, allowing more opportunity for look up triggers. Additionally, since texting agents move slower, they take longer to reach their goals so they remain in the scene for longer, which means there is more time for look ups to trigger.

6.2.3 Visual Attention Model Parameters

The results of the comparison across distraction types for refocus events, collisions, and path length is shown in Figure 6.6. For look up events, talking had the most while talking and reading were approximately the same. This is expected as talking has a much larger field of view so there is more opportunity for the reaction module to trigger look up behaviours due to the p_e and p_v probability checks. As expected, similar numbers of collisions as well as similar path lengths were recorded for normative and talking modalities, with talking having slightly higher collisions and path lengths. Surprisingly, reading had the most collisions in some cases, although the median for texters was higher than for readers indicating that texters generally experience more collisions. One possible reason readers may have had more collisions in some trials is that although the visual field is slightly more reduced for texters than readers, readers also move with higher speed which may result in more collisions. Path length was highest for readers and texters, and was slightly higher for readers in some trials than texters which may be explained by the increased collision rate in some trials for readers, which may cause them to be knocked away slightly from the optimal path. However, again the median was higher for texters than for readers, indicating texters generally have longer paths than readers.

6.2.4 Fuzzy Goal Model Parameters

The results showing the effect of the fuzzy semi-major and semi-minor lengths on pathlength is shown in Figure 6.7. Path length is highest when the maximum semi-minor axis length is high and the maximum semi-major axis length is low. In general however, path length tends to be higher when the maximum semi-minor axis length is higher. This is expected, as the semi-minor axis

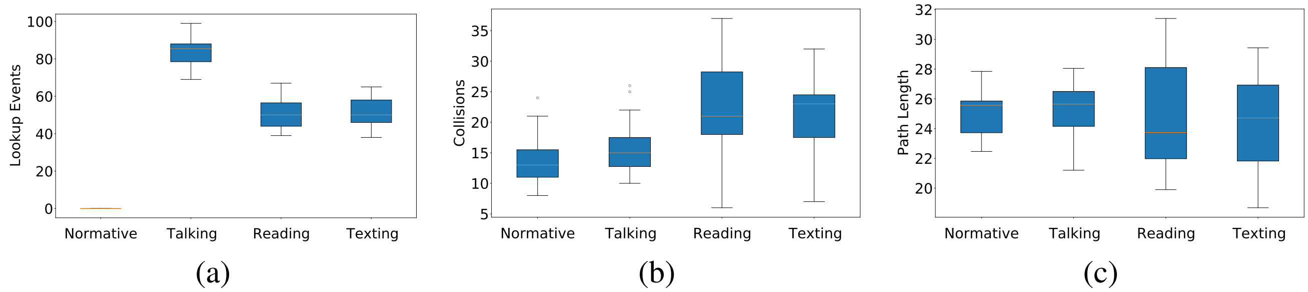

Figure 6.6: Comparisons for the different distraction models in terms of (a) look up events, (b) collisions, and (c) path length.

length defines the maximum amount of lateral deviation, which causes deviation from optimal paths and so increases path length. Higher values of semi-major axis length also tend to yield higher path lengths, but to a lesser extent and with more noise. Overshooting or undershooting may or may not result in significant deviations from the optimal path, depending on the situation. In particular, undershooting may not affect path length as much because when an agent reaches the fuzzy goal, they will continue to the next waypoint which is often in the same or similar direction they were headed. If they can't see the next waypoint, the path is recomputed, which usually results in the agent heading to the original waypoint, which would have minimal effect on the path length.

6.2.5 Qualitative Evaluation

The models were evaluated qualitatively in two indicative scenarios commonly used in crowd simulation. Two oncoming agents with obstacles is the first scenario, and the second is two crossing groups in a corridor.

Two agents in the oncoming obstacle scenario are placed at opposite sides and may potentially cross paths while simultaneously navigating around obstacles. The paths taken by the agents are shown in Fig 6.8. In all trials, both agents are performing the same distracted activity when distracted, but a different activity was used in each trial so that comparisons could be drawn between models. Normative and talking models still exhibit similar path trajectories due to the lack of lateral deviation. There are clear deviations from the optimal path for reading agents, while the trajectories for texting agents clearly show the agent that started on the left bumped the first obstacle it passed.

Additionally these evaluations were done using animated human avatars. Models were imported from Mixamo with animations, and then configured using Unity3D's Mecanim animation system.

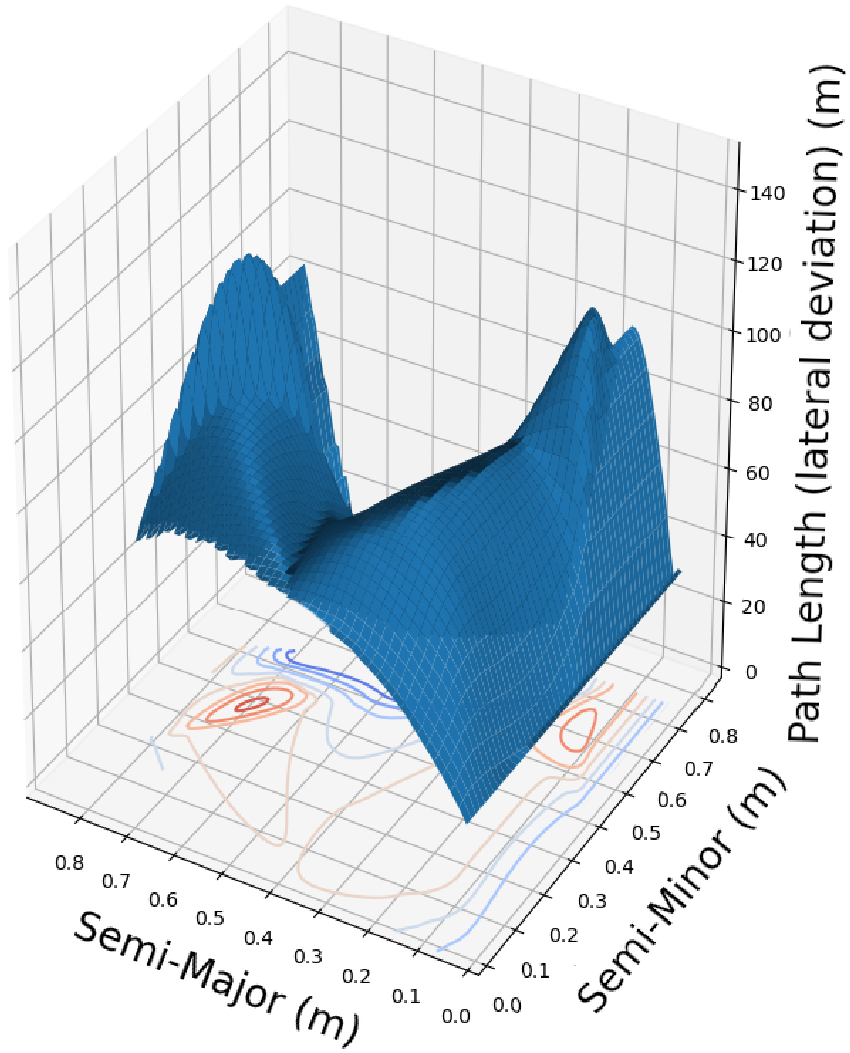

Figure 6.7: A surface representing the average agent path length as the semi major and semi minor lengths of the fuzzy goal ellipse parameters increase.

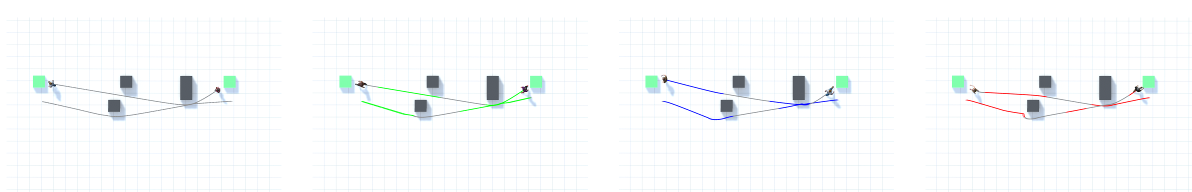

(a) (b) (c) (d)

Figure 6.8: The paths taken by (a) two normative agents, (b) two distracted agents talking on the phone, (c) two distracted agents reading, and (d) two distracted agents texting in the oncoming obstacles scenario. Green, blue, and red highlighted paths indicate the location along their path an agent was distracted.

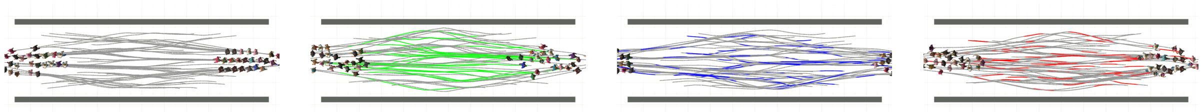

(a) (b) (c) (d)

Figure 6.9: The paths taken by (a) 50 normative agents, (b) 30 normative and 20 distracted agents talking on the phone, (c) 30 normative and 20 distracted agents reading, and (d) 30 normative and 20 distracted agents texting in the crossing groups scenario. Green, blue, and red highlighted paths indicate the location along their path an agent was distracted.

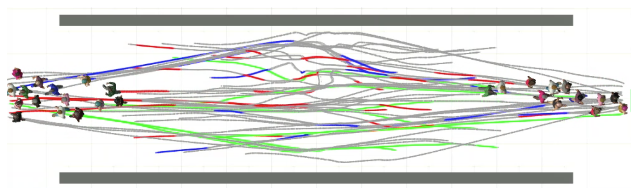

Figure 6.10: The paths taken by 30 normative and 20 distracted agents of all distraction models in the crossing groups environment. Green, blue, and red highlighted paths indicate the location along their path an agent was distracted.

In the crossing groups scenario, each group on opposite sides of the corridor had 25 agents for a total of 50 agents. Figure 6.9 shows the results for this scenario. Texting and reading distractions tended to occur more when the groups were further apart and less in the middle where the two groups met. This is due to the p_v triggering look up behaviour and also the increased number of agents in the visual field which may cause the p_e to trigger look ups. Since the talking on the phone distraction is not interrupted, this distraction type continued even in the middle where the groups met. The trajectories for agents in a trial with twenty intermittantly distracted agents performing all three activities is depicted in Figure 6.10.

Figure 6.11 shows the visual field of a distracted agent just before and right after a look up triggers. Finally, Figure 6.12 shows distracted agents crossing the street in a busy urban environment.

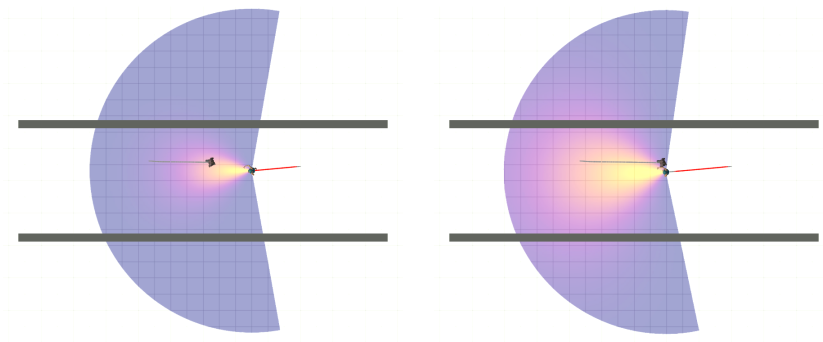

(a) (b)

Figure 6.11: A comparison of an agent's visual field (a) before and (b) after triggering a look up behaviour in the reaction module.

Figure 6.12: A screen capture of distracted and undistracted agents crossing at a crosswalk in a virtual street environment.

Chapter 7

Conclusion

Cellphone distraction while walking has become a commonplace phenomenon in many major cities, as well as a major safety concern. This is reflected in new bylaws and fines for distracted walkers in some cities, as well as the launch of safety awareness campaigns in others. The number of reported injuries where a cellphone was involved is rising every year. Crowd simulations that include these behaviours can help inform safety and urban design, and make streets safer for everyone.

This objective of this thesis was to explore whether distracted pedestrians have an impact on crowd simulations. It presents a first approach to modelling distracted behaviours in crowd simulations. The models presented use parameters based on extensive research of real world distraction studies, and therefore are representative of ground truth. In addition to parameter values, literature on real world distraction studies revealed important details that had significant impact on the design of the model, such as lateral deviation, and peripheral vision loss while focusing on a complex image constituting a high foveal load. With further research on real world distraction, I believe even more advancements could be made to the presented models.

The model presented included several modules each handling an important aspect of distraction. The fuzzy goal model was designed to simulate lateral deviation and undershoot/overshoot. The distraction module and behaviour tree was responsible for triggering distracted behaviours and setting the appropriate parameters for each behaviour in the steering layer. The visual attention module represented an agent's visual field while distracted or undistracted, and changed depending on the current activities being performed by the agent. The refocusing module determined when to trigger look up behaviour based on the active visual attention model and three Bernoulli trials conducted at every frequency check. Finally, a crowd authoring tool was designed to quickly and easily generate crowds.

Three types of distractions were tested in two sets of experiments: talking on the phone, reading on a phone, and texting. The first set of experiments used a restricted model without the visual attention and reaction modules, and with a simplified circular fuzzy goal model. Still, these experiments showed that accounting for even a few distracted pedestrians performing simple cellphone activities has a significant impant on the resulting simulations both qualitatively and quantitatively. In the Crossing Groups scenario, even a single distracted agent could have an impact on overall crowd flow and statistics depending on the activity. All distracted activities impacted crowd flow statistics after five out of fifty agents were distracted.

The second set of experiments used all the modules and an extended fuzzy goal ellipsoid model, and showed the effect each parameter had on various experiment measures. These experiments demonstrated the ability of the user to change the outcome by modifying parameters. The visual attention model, which works by convolving two Gaussians togeting in polar coordinates, was also demonstrated to have a clear impact on crowd simulation with distracted pedestrians, especially when compared with the previous experiments. Computational perfomance was also measured, showing the system in its current state is capable of simulating over 200 distracted agents without a significant increase in frame time. Some surprising results were discovered, such that reading actually causes more collisions in simulations than texting, which is likely due to the faster speed reading agents are moving at. Finally, some qualitative results showed avatars walking around with cellphones using a dynamic IK system, where the IK targets move in random directions around their center in order to create a hand bobbing effect and avoid rigid looking animations.

Future work may include direct comparisons to ground truth, and additional models for other cellphone distraction types or non-cellphone related distractions, as well as models for other non-normative gait patterns such as those seen in Cerebral Palsy. Extensions could be made to the reaction module such as look up behaviour triggered by sounds, which may not be within the agent's visual field and may or may not also be represented with egocentric fields. Diversified behaviours including group behaviours that contribute to heterogeneity would help create realistically diversified and high fidelity crowds.

Bibliography

[1] Valentina Agostini, Francesco Lo Fermo, Giuseppe Massazza, and Marco Knaflitz. Does texting while walking really affect gait in young adults? *Journal of neuroengineering and rehabilitation*, 12(1):86, 2015.

[2] Jan M Allbeck and Norman I Badler. Creating crowd variation with the ocean personality model. 2008.

[3] Alexa Angster, Michael Frank, and David Lester. An exploratory study of students' use of cell phones, texting, and social networking sites. *Psychological reports*, 107(2):402–404, 2010.

[4] Gérard Bailly, Stephan Raidt, and Frédéric Elisei. Gaze, conversational agents and face-to-face communication. *Speech Communication*, 52(6):598–612, 2010.

[5] Andrew Best, Sahil Narang, and Dinesh Manocha. Real-time reciprocal collision avoidance with elliptical agents. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 298–305. IEEE, 2016.

[6] Jaeyun Cha, Hyunjin Kim, Jaemyoung Park, and Changho Song. Effects of mobile texting and gaming on gait with obstructions under different illumination levels. *Physical Therapy Rehabilitation Science*, 4(1):32–37, 2015.

[7] Michael E Cinelli and Aftab E Patla. Locomotor avoidance behaviours during a visually guided task involving an approaching object. *Gait & posture*, 28(4):596–601, 2008.

[8] Catherine Coccia and Carol A Darling. Having the time of their life: College student stress, dating and satisfaction with life. *Stress and Health*, 32(1):28–35, 2016.

[9] NCD Risk Factor Collaboration et al. A century of trends in adult human height. *Elife*, 5:e13410, 2016.

[10] Michele Colledanchise and Petter Ögren. *Behavior trees in robotics and AI: An introduction*. CRC Press, 2018.

[11] Sean Curtis, Stephen J Guy, Basim Zafar, and Dinesh Manocha. Virtual tawaf: A case study in simulating the behavior of dense, heterogeneous crowds. In *2011 IEEE International Conference on Computer Vision Workshops (ICCV Workshops)*, pages 128–135. IEEE, 2011.

[12] Funda Durupinar, Jan Allbeck, Nuria Pelechano, and Norman Badler. Creating crowd variation with the ocean personality model. In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems-Volume 3*, pages 1217–1220. International Foundation for Autonomous Agents and Multiagent Systems, 2008.

[13] Brett R Fajen and William H Warren. Behavioral dynamics of steering, obstable avoidance, and route selection. *Journal of Experimental Psychology: Human Perception and Performance*, 29(2):343, 2003.

[14] John Funge, Xiaoyuan Tu, and Demetri Terzopoulos. Cognitive modeling: knowledge, reasoning and planning for intelligent characters. 1999.

[15] Stephen J Guy, Sean Curtis, Ming C Lin, and Dinesh Manocha. Least-effort trajectories lead to emergent crowd behaviors. *Physical review E*, 85(1):016110, 2012.

[16] Stephen J Guy, Sujeong Kim, Ming C Lin, and Dinesh Manocha. Simulating heterogeneous crowd behaviors using personality trait theory.

[17] Shigeru Haga, Ayaka Sano, Yuri Sekine, Hideka Sato, Saki Yamaguchi, and Kosuke Masuda. Effects of using a smart phone on pedestrians' attention and walking. *Procedia Manufacturing*, 3:2574–2580, 2015.

[18] M. Brandon Haworth. *Biomechanical Locomotion Heterogeneity in Synthetic Crowds*. PhD thesis, York University, Toronto, Canada, November 2019.

[19] Dirk Helbing and Peter Molnar. Social force model for pedestrian dynamics. *Physical review E*, 51(5):4282, 1995.

[20] Randall Hill. Modeling perceptual attention in virtual humans. In *Proceedings of the 8th Conference on Computer Generated Forces and Behavioral Representation*, pages 563–573. Citeseer, 1999.

[21] Ira E Hyman Jr, S Matthew Boss, Breanne M Wise, Kira E McKenzie, and Jenna M Caggiano. Did you see the unicycling clown? inattentional blindness while walking and talking on a cell phone. *Applied Cognitive Psychology*, 24(5):597–607, 2010.

[22] Ira E Hyman Jr, Benjamin A Sarb, and Breanne M Wise-Swanson. Failure to see money on a tree: inattentional blindness for objects that guided behavior. *Frontiers in psychology*, 5:356, 2014.

[23] Mitsuo Ikeda and Tetsuji Takeuchi. Influence of foveal load on the functional visual field. *Perception & Psychophysics*, 18(4):255–260, 1975.

[24] Laurent Itti, Christof Koch, and Ernst Niebur. A model of saliency-based visual attention for rapid scene analysis. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, (11):1254–1259, 1998.

[25] SoYeong Jeon, ChoRong Kim, SunHae Song, and GyuChang Lee. Changes in gait pattern during multitask using smartphones. *Work*, 53(2):241–247, 2016.

[26] Tessa Jones. Cell phone use while walking across campus: an observation and survey. *Elon Journal of Undergraduate Research in Communications*, 5(1), 2014.

[27] Eunjung Ju, Myung Geol Choi, Minji Park, Jehee Lee, Kang Hoon Lee, and Shigeo Takahashi. Morphable crowds. In *ACM Transactions on Graphics (TOG)*, volume 29, page 140. ACM, 2010.

[28] Pei-Chun Kao, Christopher I Higginson, Kelly Seymour, Morgan Kamerdze, and Jill S Higginson. Walking stability during cell phone use in healthy adults. *Gait & posture*, 41(4):947–953, 2015.

[29] Mubbasir Kapadia, Nuria Pelechano, Jan Allbeck, and Norm Badler. *Virtual Crowds: Steps Toward Behavioral Realism*. Morgan & Claypool Publishers, 2015.

[30] Ioannis Karamouzas, Peter Heil, Pascal Van Beek, and Mark H Overmars. A predictive collision avoidance model for pedestrian simulation. In *International workshop on motion in games*, pages 41–52. Springer, 2009.

[31] Ioannis Karamouzas, Brian Skinner, and Stephen J Guy. Universal power law governing pedestrian interactions. *Physical review letters*, 113(23):238701, 2014.

[32] Sonu Chopra Khullar and Norman I Badler. Where to look? automating attending behaviors of virtual human characters. *Autonomous Agents and Multi-Agent Systems*, 4(1-2):9–23, 2001.

[33] Elena Kokkinara, Oyewole Oyekoya, and Anthony Steed. Modelling selective visual attention for autonomous virtual characters. *Computer Animation and Virtual Worlds*, 22(4):361–369, 2011.

[34] Melissa Kremer, Brandon Haworth, Mubbasir Kapadia, and Petros Faloutsos. Modeling distracted agents in crowd simulations. Springer, 2020.

[35] James J Kuffner and J-C Latombe. Fast synthetic vision, memory, and learning models for virtual humans. In *Proceedings Computer Animation 1999*, pages 118–127. IEEE, 1999.

[36] Eric M Lamberg and Lisa M Muratori. Cell phones change the way we walk. *Gait & posture*, 35(4):688–690, 2012.

[37] Kang Hoon Lee, Myung Geol Choi, Qyoun Hong, and Jehee Lee. Group behavior from video: a data-driven approach to crowd simulation. In *Proceedings of the 2007 ACM SIGGRAPH/Eurographics symposium on Computer animation*, pages 109–118. Eurographics Association, 2007.

[38] HW Leibowitz and Stuart Appelle. The effect of a central task on luminance thresholds for peripherally presented stimuli. *Human Factors*, 11(4):387–391, 1969.

[39] Sammy Licence, Robynne Smith, Miranda P McGuigan, and Conrad P Earnest. Gait pattern alterations during walking, texting and walking and texting during cognitively distractive tasks while negotiating common pedestrian obstacles. *PLoS one*, 10(7):e0133281, 2015.

[40] WC Maples, Wes DeRosier, Richard Hoenes, Rodney Bendure, and Sherl Moore. The effects of cell phone use on peripheral vision. *Optometry-Journal of the American Optometric Association*, 79(1):36–42, 2008.

[41] Judith Mwakalonge, Saidi Siuhi, and Jamario White. Distracted walking: Examining the extent to pedestrian safety problems. *Journal of traffic and transportation engineering (English edition)*, 2(5):327–337, 2015.

[42] Sahil Narang, Andrew Best, Tanmay Randhavane, Ari Shapiro, and Dinesh Manocha. Pedvr: Simulating gaze-based interactions between a real user and virtual crowds. In *Proceedings of the 22nd ACM conference on virtual reality software and technology*, pages 91–100. ACM, 2016.

[43] Ulric Neisser. The control of information pickup in selective looking. *Perception and its development: A tribute to Eleanor J. Gibson*, pages 201–219, 1979.

[44] Daniel Niederer, Anke Bumann, Yvonne Mühlhauser, Mareike Schmitt, Katja Wess, Tobias Engeroff, Jan Wilke, Lutz Vogt, and Winfried Banzer. Specific smartphone usage and cognitive performance affect gait characteristics during free-living and treadmill walking. *Gait & posture*, 62:415–421, 2018.

[45] Jan Ondřej, Julien Pettré, Anne-Hélène Olivier, and Stéphane Donikian. A synthetic-vision based steering approach for crowd simulation. *ACM Transactions on Graphics (TOG)*, 29(4):1–9, 2010.

[46] Sébastien Paris, Julien Pettré, and Stéphane Donikian. Pedestrian reactive navigation for crowd simulation: a predictive approach. In *Computer Graphics Forum*, volume 26, pages 665–674. Wiley Online Library, 2007.

[47] Christopher JD Patten, Albert Kircher, Joakim Östlund, and Lena Nilsson. Using mobile telephones: cognitive workload and attention resource allocation. *Accident analysis & prevention*, 36(3):341–350, 2004.

[48] Nuria Pelechano, Kevin O'Brien, Barry Silverman, and Norman Badler. Crowd simulation incorporating agent psychological models, roles and communication. Technical report, PENNSYLVANIA UNIV PHILADELPHIA CENTER FOR HUMAN MODELING AND SIMULATION, 2005.

[49] Sam Perlmutter, David M Cades, Michelle F Heller, Robert Giachetti, Steven R Arndt, and Joseph B Sala. Effects of mobile technology use on walking. 58(1):2166–2170, 2014.

[50] Christopher Peters and Carol O'Sullivan. Synthetic vision and memory for autonomous virtual humans. In *Computer Graphics Forum*, volume 21, pages 743–752. Wiley Online Library, 2002.

[51] Sara Pizzamiglio, Usman Naeem, Hassan Abdalla, and Duncan L Turner. Neural correlates of single-and dual-task walking in the real world. *Frontiers in human neuroscience*, 11:460, 2017.

[52] Prudence Plummer, Sarah Apple, Colleen Dowd, and Eliza Keith. Texting and walking: Effect of environmental setting and task prioritization on dual-task interference in healthy young adults. *Gait & posture*, 41(1):46–51, 2015.

[53] Prudence Plummer-D'Amato, Briana Brancato, Mallory Dantowitz, Stephanie Birken, Christina Bonke, and Erin Furey. Effects of gait and cognitive task difficulty on cognitive-motor interference in aging. *Journal of aging research*, 2012, 2012.

[54] Paphawee Prupetkaew, Vipul Lugade, Teerawat Kamnardsiri, and Patima Silsupadol. Cognitive and visual demands, but not gross motor demand, of concurrent smartphone use affect laboratory and free-living gait among young and older adults. *Gait & posture*, 68:30–36, 2019.

[55] Phillip Quon, Kelcie Lahey, Mackenzie Grisdale, Brent Hagel, George Frost, Kathy Belton, and April Elliott. Prevalence of distracted walking with mobile technology: an observational study of calgary and edmonton high school students. *Canadian journal of public health*, 110(4):506–511, 2019.

[56] Craig W. Reynolds. Flocks, herds and schools: A distributed behavioral model. *SIGGRAPH Comput. Graph.*, 21(4):25–34, August 1987.

[57] Siobhan M Schabrun, Wolbert van den Hoorn, Alison Moorcroft, Cameron Greenland, and Paul W Hodges. Texting and walking: strategies for postural control and implications for safety. *PloS one*, 9(1):e84312, 2014.

[58] Daniel J Simons and Christopher F Chabris. Gorillas in our midst: Sustained inattentional blindness for dynamic events. *perception*, 28(9):1059–1074, 1999.

[59] Shawn Singh, Mubbasir Kapadia, Glenn Reinman, and Petros Faloutsos. Footstep navigation for dynamic crowds. *Computer Animation and Virtual Worlds*, 22(2-3):151–158, 2011.

[60] Jamie Snape, Jur Van Den Berg, Stephen J Guy, and Dinesh Manocha. The hybrid reciprocal velocity obstacle. *IEEE Transactions on Robotics*, 27(4):696–706, 2011.

[61] Tino Stöckel and Anett Mau-Moeller. Cognitive control processes associated with successful gait performance in dual-task walking in healthy young adults. *Psychological Research*, pages 1–11, 2019.

[62] David L Strayer, Jason M Watson, and Frank A Drews. Cognitive distraction while multitasking in the automobile. 2011.

[63] D. Thalmann and S.R. Musse. *Crowd Simulation*. Springer London, 2012.

[64] Leah L Thompson, Frederick P Rivara, Rajiv C Ayyagari, and Beth E Ebel. Impact of social and technological distraction on pedestrian crossing behaviour: an observational study. *Injury prevention*, 19(4):232–237, 2013.

[65] Yu Tian, Yifei Huang, Jibo He, and Kunlin Wei. What affects gait performance during walking while texting? a comparison of motor, visual and cognitive factors. *Ergonomics*, 61(11):1507–1518, 2018. PMID: 30009681.

[66] John K Tsotsos. Motion understanding: Task-directed attention and representations that link perception with action. *International Journal of Computer Vision*, 45(3):265–280, 2001.

[67] Jur Van Den Berg, Stephen J Guy, Ming Lin, and Dinesh Manocha. Reciprocal n-body collision avoidance. In *Robotics research*, pages 3–19. Springer, 2011.

[68] David Wolinski, Ming C Lin, and Julien Pettré. Warpdriver: context-aware probabilistic motion prediction for crowd simulation. *ACM Transactions on Graphics (TOG)*, 35(6):1–11, 2016.

[69] Zhuxin Xue, Qing Dong, Xiangtao Fan, Qingwen Jin, Hongdeng Jian, and Jian Liu. Fuzzy logic-based model that incorporates personality traits for heterogeneous pedestrians. *Symmetry*, 9(10):239, 2017.

[70] Syuji Yoshiki, Hiroshi Tatsumi, Kayoko Tsutsumi, Toru Miyazaki, and Takuya Fujiki. Effects of smartphone use on behavior while walking. *Urban and Regional Planning Review*, 4:138–150, 2017.

[71] Kyung-Hoon Yu, Jae-Hun Shim, Sung-Dae Choung, and Hye-Seon Jeon. Effect of using a cell phone on gait parameters in healthy young adults: Texting and texting while listening to music. *Journal of Korean Society of Physical Medicine*, 10(4):25–31, 2015.

[72] Qinxin Yu and Demetri Terzopoulos. A decision network framework for the behavioral animation of virtual humans. In *Proceedings of the 2007 ACM SIGGRAPH/Eurographics symposium on Computer animation*, pages 119–128. Eurographics Association, 2007.

[73] Mohamed H Zaki and Tarek Sayed. Exploring walking gait features for the automated recognition of distracted pedestrians. *IET Intelligent Transport Systems*, 10(2):106–113, 2016.

[74] Liping Zheng, Dang Qin, Yajun Cheng, Lin Wang, and Lin Li. Simulating heterogeneous crowds from a physiological perspective. *Neurocomputing*, 172:180–188, 2016.

[75] George Kingsley Zipf. Human behavior and the principle of least effort. 1949.