Modelling behavior in agent based simulations of crowds

Vaisagh Viswanathan

2015

Vaisagh Viswanathan. (2015). Modelling behavior in agent based simulations of crowds.
Doctoral thesis, Nanyang Technological University, Singapore.

<https://hdl.handle.net/10356/62668>

<https://doi.org/10.32657/10356/62668>

**MODELLING BEHAVIOR IN AGENT BASED
SIMULATIONS OF CROWDS**

VAISAGH VISWANATHAN

School of Computer Engineering

A thesis submitted to the Nanyang Technological University
in partial fulfilment of the requirement for the degree of

Doctor of Philosophy

2015

Acknowledgements

I would like to express my sincere gratitude to my supervisor, **Dr. Bo An** for his incredible understanding, help and support in a very difficult situation. I would also like to thank my co-supervisor and mentor, **Dr. Michael Lees** for giving me the opportunity and the inspiration to pursue my dreams. Right from the days of co-supervising my Final Year Project, he has always been a reliable and uncompromising critic who's been there to help, advise and discuss about everything ranging from my research to teaching to technology and life in general. The weekly team meetings over coffee and, more recently, over hangouts, and his emphasis on encouraging everyone in his team to discuss their projects with each other not only helped me develop a lot of my ideas but made PhD life a lot more interesting when things got boring; and bearable when the going got tough.

Prof. Peter Slood has my sincerest gratitude for taking time out of his busy schedule to give me valuable feedback, advise and support for my research. I would also like to thank **Dr. Cheong Siew Ann** for giving me the chance to work with and learn from him. A special thanks goes out to **Prof. Cai Wentong** for introducing me to Crowd Simulation and giving me the chance to be a part of several interesting projects over the past several years.

I'm grateful to **Dr. Heiko Aydt** and **Dr. Hu Nan** for not only letting me be a part of some of the really interesting work they were doing but, more importantly, for sharing their experience in working on a PhD at NTU which helped me avoid and handle most of the major and minor hurdles that I faced as a PhD Candidate. I must also acknowledge **Debraj Roy** and **Nasri Bin Othman** for being sounding boards for my ideas and friends in an environment where I've sometimes felt lost. My thesis would not have been complete without the invaluable help of **Cheong Eu Lee** in making sense of the numbers.

I would also like to express my sincere gratitude to the School of Computer Engineering, NTU and the Parallel and Distributed Computing Center for providing me with the opportunity, the funds and, above all, the excellent facilities to pursue a PhD. A special thanks is also due to **Irene** (Ng-Goh Siew Lai), **Kenneth** (Poliran Kenneth Caballes) and **Saen** (Lok Chun Saen) for making working at PDCC so easy and convenient.

Finally, I would like to thank my family for their unflinching support and belief in me; and **Aparna** for listening to my constant complaints and rants, and always being there to cheer me up.

Contents

1	Introduction	1
1.1	Problem Statement	4
1.2	Key Contributions and Scope of this Thesis	5
1.3	Organization of the Thesis	6
2	Literature Review	7
2.1	Multidisciplinary Nature	8
2.2	Current Understanding of Human Behavior During Egress	10
2.2.1	Pre-evacuation behavior	10
2.2.2	Crowd behavior	11
2.2.3	On the rationality of crowds	15
2.2.4	Summary	16
2.3	Computational Models of Egress	17
2.3.1	Components of a computational model of egress	18
2.3.1.1	A model of the environment	18
2.3.1.2	A model of the individuals engaging in egress	18
2.3.2	An overview of existing computational models	20
2.3.2.1	Network or Queuing Theory based approaches	20
2.3.2.2	Flow models	21
2.3.2.3	Environment Control Based Models	21
2.3.2.3.1	Cellular Automaton based approaches:	21
2.3.2.3.2	Lattice gas models:	22
2.3.2.3.3	Floor field models:	22
2.3.2.4	Agent based approaches	23
2.3.2.5	Hybrid approaches	24
2.3.3	Detailed discussion of significant computational models	24
2.3.3.1	Pires's model of pre-evacuation behavior	24
2.3.3.2	Legion	24
2.3.3.3	MACES + PMFServ	25
2.3.3.4	Collective Panic Behavior Model	26
2.3.3.5	Exodus	28
2.3.3.6	ESCAPES	30

2.3.3.7	MASSEgress	31
2.3.4	Summary	32
2.4	Summary of Literature Review	33
3	The Building Blocks of a Behavior Model for Egress Simulation	35
3.1	The Building Blocks of the Egress Process	36
3.1.1	Perception	36
3.1.2	Event identification	37
3.1.3	Spatial knowledge utilization	37
3.1.4	Navigation	38
3.2	Shortcomings in Existing Models	41
3.2.1	Perception: the lack of a realistic approach	41
3.2.2	Event identification: modelling pre-evacuation behavior	41
3.2.3	Knowledge modelling: the effect of partial knowledge	41
3.2.4	Navigation: analyzing motion planning systems	42
3.3	Contributions of the Remaining Chapters	42
3.4	Summary	43
4	An Information Processing Based Approach to Modelling Perception	45
4.1	Limits of Human Perception	47
4.2	A Brief Introduction to Motion Planning Systems	48
4.3	The Information Based Perception Model	51
4.3.1	Clustering	51
4.3.2	Sensing	54
4.3.3	Filtering	55
4.4	Model Evaluation	57
4.4.1	Model demonstration	57
4.4.1.1	Group Based Perception	58
4.4.1.2	Effects of multi layered clustering	60
4.4.1.3	Filtering necessitates Group Based Perception	62
4.4.1.4	Effect of filtering of percept information	62
4.4.2	Validation against standard scenarios	64
4.4.2.1	Corridor (1 <i>mma</i>)	65
4.4.2.2	Intersection (2 <i>mma</i>)	66
4.4.2.3	Circle (multiple <i>mma</i>)	66
4.4.3	Experimental validation	68
4.4.3.1	Experiment setup	69
4.4.3.2	Quantitative measurements of experimental data	71
4.4.3.3	Results and comparison	73
4.5	Summary	75

5	Modelling Pre-evacuation Behavior	77
5.1	Related Work	78
5.1.1	On modelling pre-evacuation behavior	78
5.1.2	Existing models	80
5.2	A Bucket Model of Event Identification	82
5.3	Results	84
5.3.1	Experiment 1: the effect of fire alarm ambiguity	90
5.3.2	Experiment 2: informed and informing agents	91
5.3.3	Experiment 3: is milling important?	94
5.3.4	Experiment 4: the importance of spatial knowledge propagation	96
5.4	Summary	97
6	A Game-based Investigation into the Role of Memory in Human Exploration and Navigation	99
6.1	Literature Review	101
6.1.1	Working memory	101
6.1.2	The building blocks of spatial knowledge	102
6.1.3	Wayfinding in indoor environments	103
6.2	Benefits of a Game Based Approach	106
6.3	Setup of the Experiment	107
6.3.1	The game	107
6.3.2	The Minecraft gaming environment	110
6.3.3	Participants	110
6.4	Analysis	111
6.4.1	Types of exploring agents	111
6.4.2	Markov Agents	113
6.4.2.1	Calculation of Markov data	114
6.5	Results	117
6.5.1	Room visit frequencies	117
6.5.2	Expected coverage given number of hops	121
6.5.3	Expected hops given coverage	122
6.5.4	Empirical analysis	122
6.5.4.1	Existence of decision points	124
6.5.4.2	Location recognition and memory	125
6.6	Summary	127
7	Quantitative Comparison Between Motion Planning Systems for Evacuation Planning and Evaluation	129
7.1	Motion Planning Models	131
7.1.1	Lattice gas model	131
7.1.2	Social force model	133

7.1.3	Reciprocal velocity obstacles model	134
7.2	Methodology	135
7.2.1	Measurements	135
7.2.2	DISTATIS	137
7.3	Results and Discussion	138
7.3.1	Evacuation time	140
7.3.2	Passage density	142
7.3.3	Total distance traveled distribution and inconvenience	143
7.3.4	Evacuation rate	144
7.3.5	DISTATIS comparison	146
7.4	Summary	146
8	Conclusion and Future Work	149
8.1	Modelling Human Cognitive Limits in Perception	151
8.2	A Novel Model for Event Identification	152
8.3	A Game Based Analysis of Indoor Wayfinding Behavior	153
8.4	A Comparison of Existing Motion Planning Systems	155
8.5	Future Work	156
8.6	Conclusion	157
A	Comparison of Information Based Perception against Controlled Experiments	159
B	Validity Check for Markov analysis	161
B.1	Effect of Dataset Size	161
B.2	Decision Base Size at Decision Points	161

List of Figures

2.1	Strategies for egress	14
2.2	EXODUS model	29
3.1	The process of egress	36
3.2	Detailed navigation model	39
4.1	Navigation architecture	49
4.2	Perception and acting	51
4.3	Breakdown of the perception process	52
4.4	A multi-layered clustering approach	54
4.5	Distance score	56
4.6	Angle score	57
4.7	Experiment 1: Group Based Perception	58
4.8	Experiment 2: Effect of multi layered clustering	60
4.9	Experiment 2: Quantitative analysis	61
4.10	Experiment 3: The necessity of Group Based Perception	63
4.11	Experiment 4: Effect of filtering of percept information	64
4.12	Corridor experiment screenshots	65
4.13	Intersection experiment screenshots	67
4.14	Circle experiment screenshots	68
4.15	Schematic of layout of experiment area [1].	69
4.16	Actual setup of Hu's experiments [1]	70
4.17	The data extraction process.	71
4.18	An example of the video extraction process	72
4.19	The trajectories of participants can be grouped into 3 categories	72
4.20	Experiment evaluation metrics	73
4.21	Layout of experiment in simulation	73
4.22	Maximum deviation measurements from simulation	75
4.23	Inter agent distances over time.	76
5.1	An illustration of the bucket model	83
5.2	The Environment Layout	85
5.3	Survival rate heat map for immediate evacuation	87

5.4	Survival rate heat map for delayed evacuation	88
5.5	Difference between immediate and delayed evacuation	89
5.6	Experiment 1: The effect of fire alarm ambiguity	91
5.7	Experiment 2: The effect of having informed occupants	93
5.8	Experiment 2: The effect of having informed occupants (contd.)	94
5.9	Experiment 3: The importance of milling	95
5.10	Experiment 4: The importance of spatial knowledge propagation	97
6.1	Floor plans of the three floors	108
6.2	Picture gallery	109
6.3	Graphical representation of floor plan	112
6.4	Markov data explanation	115
6.5	Visualization of room visit frequencies	118
6.6	Room visit frequencies normalized to floor	120
6.7	Expected coverage after 253 hops	121
6.8	Minimum expected number of hops for given coverage	123
6.9	Existence of definite decision points	124
6.10	Layout of dorm corridor	125
6.11	Player behavior at junction under consideration	126
6.12	Player behavior at dorm corridor	127
7.1	Working of lattice gas model	132
7.2	The environment setup used for the simulations in this chapter. . . .	136
7.3	Zonal division of environment	137
7.4	Visualization of DISTATIS workflow	139
7.5	Evacuation time distributions	140
7.6	Zoned evacuation time distributions	141
7.7	Passage density maps	142
7.8	Distributions of total distance traveled	143
7.9	Distributions of inconveniences	144
7.10	Visualization of flow rates	145
7.11	Flow rates during Sichuan earthquake equation	146
7.12	DISTATIS analysis of three models	147
A.1	Speed plot comparison	160
B.1	Dataset size dependence of calculations	162
B.2	Comparison of the decision base sizes.	162

List of Tables

3.1	The building blocks of human behavior during egress	40
4.1	Quantitative analysis of Group Based Perception	59
4.2	Table summarizing effort expended for standard scenarios	69
4.3	Real world maximum deviation [1]	74
4.4	Maximum deviation measurements from simulations	74
5.1	Factors affecting Evacuation Behavior.	81
6.1	Summary of symbols and their meaning.	114
A.1	Real world trajectory lengths [1]	159
A.2	Simulation trajectory lengths	160

Summary

Crowd simulation is gaining an increasing amount of importance in recent years for a variety of purposes ranging from movies and games to safe design of buildings, planning for major events like the Olympics and for preparing for emergencies. Over the last few decades, crowd simulation models have come a long way from the simpler network based approaches to simulations using agent based modelling (ABM) which have a much greater amount of detail. The bottom up approach of ABM allows modelers to consider the complexity of human behavior in much more detail than was traditionally feasible. It is possible to consider thousands of individuals with their individual behavior characteristics occupying and interacting in complex indoor environments like shopping malls.

However, existing computational models fail to take into consideration the large body of work on human behavior during emergencies and make unsubstantiated assumptions like perfect knowledge of the layout and immediate evacuation on hearing a fire alarm. In this thesis, the existing work on human behavior during emergency egress is studied and perception, event identification, spatial knowledge acquisition and navigation are identified as the four key building blocks of a behavior model for agent based crowd simulation. Following this, key shortcomings in existing models of each of these are identified and addressed.

Existing simulations model perception as a simple process of visual sensing of a fixed circular range. A more realistic perception model considering the limitations of human information processing capacity is proposed and its usefulness is demonstrated through its effect on existing motion planning systems. Pre-evacuation behavior is generally either simplified or ignored despite studies showing its importance. A model for simulating pre-evacuation behavior is introduced and its importance in developing better strategies for egress management is demonstrated. Due to limited understanding of the impact of partial spatial knowledge on indoor wayfinding, existing models assume complete knowledge. A game-based methodology was used to reveal patterns in indoor wayfinding like decision points and the impact of short-term memory which enables more accurate modeling of evacuation. The effect that movement model choice has on egress dynamics has never been studied systematically. The final contribution is a method for quantitatively comparing motion planning systems and a comparison of three popular motion planning systems that revealed the importance of zoned evacuation time as an important metric for validation.

Chapter 1

Introduction

At about 3.15 pm on December 30, 1903, the full capacity crowd at the Iroquois theater in Chicago were just getting prepared for the second act of *Mr. Bluebeard*, when tragedy struck. An arc light shorted out and sparks ignited the curtain and started a fire that quickly engulfed the building. 605 deaths were reported in what is still the deadliest single building fire in United States history. Subsequent analysis revealed that the death toll would have been much less had the building design been less confusing. One of the most analyzed among recent such tragedies is the Love Parade disaster of 2010 which occurred in Duisburg, Germany. A crowd rush in the festival area lead to 21 deaths and injuries to more than 500 people. Helbing and Mukerji [2] concluded that the major reason for the death and destruction was the inadequate capacity of the holding area and the fact that the parade area could only be entered and exited through one tunnel. More than a century has passed since the Iroquois Theater Fire, but stampedes and deaths from people trying to escape from enclosed spaces are still very common and still seem to be caused by inadequacies in planning for emergencies.

Several of these deaths during egress are preventable with proper planning and careful construction of buildings. Modelling and simulation is one of the primary tools through which planning is done these days. For example, the Sydney Olympics made use of the crowd simulation software, LEGION [3] to test the facilities for their ability to accommodate crowds and handle emergency evacuations. In fact, crowd simulation is an area that has been the subject of a significant amount of multi-disciplinary work over the last few decades [3–5]. Its applications range from simulating crowds for movies [6, 7] and games [8, 9] to analyzing pedestrian behavior [10–12] and preparing for fire evacuations and similar emergencies [13–15]. Animation of large crowds in the award winning movie, *The Lord of the Rings* was done using a commercial crowd simulation software called MASSIVE [6]. Despite the range of uses, some of the most frequent users of crowd simulation systems are civil defense authorities, who make use of these simulations to study, evaluate and formulate strategies for controlling crowds and for tackling emergencies that can emerge, like in the case of the

Sydney Olympics. While modelling crowds, or more specifically, crowds engaging in egress, the structure of the building, the spread of fire or smoke and their effect and, most importantly, the behavior and reactions of the evacuees need to be considered. This makes it an inherently difficult system to simulate.

Recently, in various fields, there has been an increasing emphasis on the importance of recognizing the *complexity* of real world systems [16]. Effort is being put into modelling these complexities to better understand and predict what happens in the real world. One of the most popular methodologies for tackling this complexity has been through the use of Agent Based Modelling (ABM). An Agent Based Model is made up of multiple heterogeneous intelligent entities known as agents interacting in an environment. The bottom up approach of modelling each agent's behavior often allows the modeler to implement existing theories of individual behavior directly without having to abstract away too many details. This approach is useful because it is easier to study and formulate theories at an individual level; apply this theory to model a real world system populated by many such individuals; and finally simulate it. Higher level patterns that emerge from this microlevel modelling can be analyzed and studied to learn more about the system and make predictions. Thus agent based models are ideal for models where there is enough information about the behavior of individual interacting entities.

A crowd engaging in egress, with its complex system of several thinking, feeling people, a building structure that shapes their movements and other entities like fires, smoke and fire alarms, is one such real world system where the complexity of the situation needs to be considered but is often ignored for simplicity. Managing egress is one of the key challenges faced during all kinds of disasters ranging from fires to terrorist attacks to natural disasters like hurricanes and tsunamis that seem to be becoming increasingly common today. The first step towards managing such situations is understanding how people behave when confronted with these situations. Several studies in psychology and sociology over the past decades have helped us understand a lot more about human behavior during egress than we did even a few years ago. Perhaps, the most interesting (and surprising) among these, is the fact that irrational *mass panic*, that is found in news reports of incidents, does not seem to occur and definitely does not spread through a crowd [17–22]. The general consensus in the scientific community today is that crowds do not start behaving irrationally when in such situations and generally try their best to help others. Most, if not all, people behave rationally within the bounds of what knowledge they have of the situation in a phenomenon referred to as *bounded rationality*. The assumption of irrational panic is just one among several reasons for many psychologists and sociologists remaining unconvinced about the efficacy and accuracy of the results produced by today's simulations [20, 23, 24]. Several details of human behavior are abstracted away for simplicity, when in fact they play a key role in determining the egress dynamics.

Studies emphasizing the importance of bounded rationality and the post-hoc analysis of disasters like those in Iroquois theater and the Love Parade seem to emphasize the important role played by information and the human brain’s ability to process information. The limited capacity of the human brain to process information [25] at any given time can have a significant effect on the way humans perceive the environment and form their knowledge of the current situation. In fact, studies [26, 27] have shown that this capacity is reduced further by circumstances like stress, time constraints or old age. This can explain the apparent irrational behavior that is observed in high stress situations like emergency egress. The primary cause of deaths in most crowd disasters, as will be discussed in the subsequent chapters, is the lack of *information*, either about the current situation or the best routes for escape or some other knowledge that might be crucial for safe egress.

Models need to consider all aspects and all phases of egress, from the initial recognition of an emergency all the way to exiting to safety. The research in this thesis takes a holistic view of the entire process and identifies the deficiencies in current modelling and understanding. For example, occupants of a building rarely start evacuating at the sound of the fire alarm. The first reaction of most people is either to ignore the alarm or, at best, try to finish what they are currently doing so that they can investigate and ascertain the need to evacuate. During this investigation period, occupants try to gather more information either by exploring or communicating with other occupants to learn more about the situation. It is only once this investigation confirms the need to evacuate, that actual evacuation starts. There have been incidents like the Boland Hall Fire [28] where this pre-evacuation delay has caused the deaths of occupants. However, except in rare cases like the ESCAPES model [29], pre-evacuation behavior is never modeled. Even in these cases, the model is too simple to analyse the effect of pre-evacuation behavior in much detail.

Once evacuations starts, most models assume that people take some form of the shortest path to get to the exit. However, it may be that not all occupants have complete knowledge of the building and so may not know the best route for safe escape. Limitations of experience and human cognitive capacity mean that, unless special effort is taken, most people will end up taking inefficient routes. In rare cases that partial spatial knowledge is considered [30], people are assumed to have eidetic memory, never forgetting a room that is perceived once. It is also common to model agents that do a breadth first search in case there is no knowledge, despite there being a lack of evidence indicating that this is how humans explore unknown indoor environments.

Finally, once a path is chosen, the computational model has to simulate how the evacuee takes the path towards that exit. There are several models of *motion planning* that take very different approaches to simulating this evacuee movement. This, however, makes it very difficult to choose the right model since there does not

exist a systematic method for comparing existing ones. In fact, traditionally, motion planning models are only validated against empirical observations of pedestrian movement. Moreover, existing models generally consider motion planning as simply the process of effectively avoiding collisions with other agents that are sensed within a circular or, at best, elliptical sensor range of the agent. However, the limited information processing capacity and the cognitive limits of a human being can have an effect on how collisions are avoided and how people move. Modelling this can help produce more realistic movement behavior from existing models of navigation.

In this thesis, we aim to demonstrate the importance and the usefulness of taking into consideration the inefficiencies in a human's ability to process information in simulations of crowds and, in the process, advance the state of the art in the computational modelling of egress.

1.1 Problem Statement

Computational models of egress play a key role in furthering our understanding of egress behavior and managing crowds by preventing dangerous situations from emerging. Models and simulations need to be developed that are based on the current understanding of human behavior.

Studies over the past few decades have changed our understanding of human behavior during fire evacuations. However, existing models do not take into consideration the considerable work done in other fields [23]. Several studies [17–22] have shown how humans always behave rationally with the limited information that they have and that humans hardly ever panic and behave irrationally. However, it is still common to model panic/ irrational behavior in evacuating crowds. Pre-evacuation behavior and the search for information are two general characteristics of the fire evacuation process that are very rarely considered in egress simulation models. The few computational models of egress that do consider pre-evacuation behavior [31, 32] are rather simple. In other cases, like the models of human motion and exploration used in existing models, there are several different approaches that have been used but a consistent methodology for their evaluation and validation is lacking. Similarly, certain aspects of egress behavior, like the effect of partial knowledge, have not been studied in much detail. It is necessary to first identify the key components of a behavior model of humans during egress to improve existing models. Following this, the effect of human cognitive limits on these components can be identified and modeled and their effect analyzed.

This thesis aims at advancing the state of the art in computational modelling of egress by developing models and simulations that are based on modern theory and, where possible, experiments that were conducted. This is done by first identifying the key components of a behavior model of egress and subsequently

proposing models to identify, simulate and analyze the effect of human cognitive limits on these components. This will help produce more realistic models that help plan for and handle emergency egress better.

1.2 Key Contributions and Scope of this Thesis

The key contributions of this thesis are:

- **A multidisciplinary survey and analysis of current literature on fire evacuation and crowd behavior.** This serves two important purposes. Firstly, the multi-disciplinary survey of human behavior research provides a clear idea of our current understanding of human behavior during egress. It also provides a context for the subsequent analysis of existing computational models of egress behavior; this analysis gives a better picture of the shortcomings of existing models in taking into consideration the effects of human cognitive limits on evacuation behavior and motivates the rest of the thesis.
- **A four component breakdown of egress behavior modelling.** The reviewed literature is used as a basis for explaining egress behavior as being a result of four components: Perception, Event Identification, Spatial Knowledge Utilization and Navigation. This breakdown helps identify research problems in each of these components.
- **A novel information based perception system that can model the complexities and limitations of the human perception system.** Existing models tend to model perception as a simple process of visual sensing where the agent senses other agents in a fixed circular range around it. The proposed model takes into consideration the limitations of our information processing capacity and demonstrates the usefulness of modelling it.
- **A new model for event identification which can be used for simulating pre-evacuation behavior and the process of information seeking.** Studies have repeatedly shown pre-evacuation delay among building occupants who are unconvinced about the need to evacuate. However, this has never been computationally modeled in any detail. Besides proposing a new model for simulating pre-evacuation information seeking, the importance of modelling pre-evacuation behavior is also demonstrated.
- **A human computation based investigation of the nature of spatial memory and the role of memory in exploration of multi-storey buildings.** A key issue in modelling the effect of human cognitive limits on spatial knowledge utilization is our limited understanding. Existing methods in experimental psychology are difficult to scale; this makes it difficult to draw general

conclusions of human spatial memory. The game based investigation introduced in this thesis provides a scalable alternative to existing methods and reveals interesting insights into the role of memory in indoor wayfinding.

- **A detailed quantitative comparison between popular models for motion planning.** Movement modelling is a key component of any crowd simulation model. However, a primary difficulty in comparing existing models and determining the usefulness of a new method is a lack of a standard method for comparison and analysis. In this thesis, a quantitative comparison is performed between three of today's most popular motion planning models and a key metric for differentiating and analyzing models is identified.

1.3 Organization of the Thesis

This thesis is organized as follows: Chapter 2 provides a comprehensive review of relevant theories and existing models and also provides a critical analysis of some of the most relevant ones. The following chapter uses the literature reviewed in Chapter 2 to identify the key building blocks of human behavior during evacuation. The following chapters identify and solve issues in existing computational models of each of these building blocks. Chapter 4 presents an Information processing Based Perception Model and demonstrates and validates its capabilities and working through several experiments. Chapter 5 then presents a new model for event identification which can be used for simulating pre-evacuation behavior and also demonstrates through simulations the importance of modelling pre-evacuation behavior. Chapter 6 introduces a novel game-based experimental approach to understanding the impact that memory and non-randomness have in the human exploration process. The analysis from this game is used to study exploration of indoor environments. Finally, Chapter 7 addresses one of the key issues in movement modelling in crowds by presenting a detailed methodology for comparison of existing motion planning systems and further providing an analysis of some of the most popular motion planning models used today. Chapter 8 concludes the thesis.

Chapter 2

Literature Review

Crowds are typically seen as a large collection of individuals sharing a common location and undergoing some common experience [33]. Depending on the experience being shared, individuals in a crowd can think, behave and react in different ways. In this thesis, the focus is on crowds involved in emergency egress. *Egress* refers to the process of leaving a place, and emergency egress situations are situations in which a crowd tries to escape from danger due to fires, bomb blasts or other similar events.

Crowd modelling and egress simulation are not *new* areas of research in any sense of the word. A lot of researchers have worked on modelling and studying crowds and evacuations over the past few decades. These include psychologists, sociologists, computer engineers, civil engineers, etc. In fact, fire safety science is itself an active area of research with dedicated journals. In this chapter, the current state of research in the understanding and modelling of crowds and, more specifically, emergency egress simulation is introduced. A holistic approach is taken towards understanding and modelling the egress process through the examination of literature from all related fields. The objective of this holistic analysis is to determine the core components of the egress process and discover the steps in that process that can be investigated in this thesis.

Section 2.1 talks about the essentially inter-disciplinary nature of crowd simulation and the importance of understanding and utilizing the different viewpoints and approaches. Section 2.2 introduces the reader to some of the psychological and sociological theories of crowds and people's reaction to fires. This helps determine the core processes that need to be considered in egress modelling and simulation. Subsequently, Section 2.3 explores the landscape of computational models of egress that have been developed over the years with a special emphasis on their ability to model the processes discussed in Section 2.2. This helps motivate the rest of the thesis.

2.1 Multidisciplinary Nature

A complex system can be defined as a system with many interacting entities whose properties are well defined. Inequalities in spatial and temporal interactions at the microscopic level can lead to emergent macroscopic phenomena [34]. An emergency egress simulation is just such a complex system with a lot of interacting elements. This system is made of a large number of human beings, the building, fire, smoke, fire alarms and other entities which might be specific to the location or type of emergency. Besides physicists studying the physics behind the development and spread of fire and smoke, there are other scientists and engineers who work on computationally modelling the spread of fire or smoke and the different complexities involved. There are also fire safety experts who work on improving the effectiveness of fire alarms and PA systems and, in general, on improving the efficiency of egress. A lot of researchers also work on the element at the core of this complex system: crowds or, more generally, human beings.

Many psychologists have been studying crowds and how individuals behave in crowds for over a century now and many theories have been developed. A lot of these theories contradict each other and there is generally little consensus on which is the best. Torres [20] provides an excellent comparison of the different existing theories of crowd behavior. He does so by comparing their effectiveness in explaining the fire that occurred at the Station night club in West Warwick, United States on Feb 20, 2003. Section 2.2, which presents and explains some of the existing crowd behavior theories, uses the findings from this thesis extensively.

How people interact with each other and act in a crowd is just one aspect of the person's behavior during a fire. For instance, people don't start evacuating the building as soon as they hear an alarm or as soon as they hear there is a fire. Each *cue* that he/she observes has a certain impact on the person depending on his/her identity, social role and the circumstances. There are a lot of studies [20, 22, 26, 31, 35] that examine the effect of different cues on egress behavior and the other aspects of what we refer to as *pre-evacuation behavior*. Kuligowski [35] presented a list of various cues and how each of them affects a person. She emphasized how cues need to be perceived and interpreted and how the decision for an action is taken based on these cues.

How decisions are made is itself an active area of research. Besides the theories on crowd behavior, some studies [26, 31], also propose different ways in which humans make decisions during emergencies. There are also several other dimensions to understanding the behavior of individuals engaged in egress. Various studies [17, 36, 37] discuss the effects of culture on egress and others [38, 39] propose detailed models of emotion.

Computational models of egress try to model and simulate a crowd engaging in emergency egress so that these situations can be effectively tackled. There are several different computational models of egress and each of them have their own

advantages. The major difference in opinion among computational model developers is on the amount of abstraction to be used. Some models like network based models and flow models, do not explicitly consider humans but rather view crowds as a homogeneous entity. These approaches are simpler for the modeler and computationally less resource intensive. Despite this, some crucial details like flow rate through doorways and the rate of survival can be approximately measured and analyzed. Agent Based Models, on the other hand, use significantly lesser abstraction in modelling crowds and are thus capable of simulating the heterogeneity in crowds. These approaches take a bottom up approach; by specifying each individual's behavior it can more accurately approximate real life behavior. There are also more detailed agent based models of egress simulation that model human memory and emotions. These different modelling techniques and models and their strengths and weaknesses will be discussed in more detail in Section 2.3.

Some studies [20, 23, 24] criticize computational models made by engineers and computer scientists for not taking into account the considerable advancements in the understanding of human behavior made by psychologists and sociologists. Recently, some models [40] have made a start in this direction by recognizing the need for behavioral models. However, despite being about two decades old, Sime's quote [24]:

To date the significance of human cognition, decision making and social behavior seems to be recognized, but has not been fully incorporated into the prototype working models ...

on behavior models is still very much applicable as will be illustrated in the following pages.

For example, the idea of panic or *non-adaptive behavior* is at the center of several behavioral models [40, 41], but studies have shown that this concept is neither clearly explained [20] nor well accepted [21, 24, 37, 42–44]. Also, it is difficult to find computational models that model factors like *pre-evacuation behavior* and *incomplete knowledge* of evacuees.

As our knowledge of human behavior, emotion, social interactions and decision making expands, a more complex picture of human behavior has been emerging. The use of this knowledge in egress modelling has been limited. With current advancements in technology, there are fewer reasons to make unnecessary abstractions; thus, higher fidelity models of egress can be developed. With this purpose, the next section gives an overview of the current state of knowledge about human behavior in crowds and emergency situations.

2.2 Current Understanding of Human Behavior During Egress

As discussed earlier, a fire evacuation is a complex situation to model and simulate. The main reason for this complexity is that the system has a complicated entity at its center: a thinking, feeling and socializing human being. To accurately simulate the evacuation of a building in an emergency situation, it is necessary to understand the behavior and decision making of the people taking part in it. As briefly mentioned in Section 2.1, there are a lot of conflicting views and theories on how humans behave in emergencies and why they behave as they do. However, there are certain behaviors that are fundamental to human nature and are generally accepted as fact like the tendency to search for familiar surroundings [21, 37, 44, 45] and the constant search for information [22, 26, 46, 47]. In this section, the reader is introduced to some of the major accepted ideas on the nature of human behavior during egress and some of the major theories on the same.

2.2.1 Pre-evacuation behavior

In an ideal scenario, all occupants on hearing a fire alarm, use the nearest safe exit to leave the building. However, this is hardly the norm. In most cases, occupants are used to hearing false alarms and often do not start to evacuate until they are completely sure that it is needed. On January 19, 2000, a fire in Boland Hall in Seton Hall University in United States killed three students because they had ignored the fire alarms assuming they were false alarms [28]. This uncertainty about the authenticity of a first sign of danger isn't an isolated incident [24, 43, 44, 46–50]. So, when studying the behavior of evacuees, it is necessary to study and understand their actions right from the point the fire started [47] rather than the time at which they start evacuating.

During an emergency there are some changes in the environment that indicate that something is wrong or different from normal. These changes are called *cues* [22]. Broadly, cues can be categorized as *ambiguous* and *unambiguous* cues depending on the clarity of the signal sent to the evacuee. For example, hearing noises is ambiguous, while seeing smoke or fire is unambiguous. Cues and their effects will be discussed in more detail in Section 5.1. The following paragraphs give a brief summary of the discussion to give a broad idea of pre-evacuation behavior.

Several studies [44, 47, 51] have discussed the effect of ambiguity of cues. Most importantly, cues have to persist for a period of time before evacuees even initiate investigation which can eventually lead to evacuation if more information is found. There are several factors that influence the effect of the cue on an evacuee. This includes the evacuees intrinsic factors like age and gender. In a classic study [52], participants were placed in a room with smoke for a period of time in two different

situations. The experiment showed that when alone, 75 percent of the subjects reported the smoke. In the presence of two non-reacting others only 10 percent of the subjects reported the smoke during the experimental period. People react differently to fires depending on their familiarity with the environment and preparedness for the emergency [21, 37, 43, 46, 47, 53].

In general, the term *pre-evacuation* refers to the period of time that elapses between the start of the fire alarm (which is generally the first cue) and the person starting egress. According to several studies this process can itself be broken down into three sub-processes:

1. Perception: The fact that a cue exists does not mean that the evacuee has perceived it. For example, he could be asleep.
2. Investigation: Once some cues are perceived, evacuees start an investigation process that may consist of milling or simply exploring (Section 2.2.2).

It's only once investigation reveals sufficient information to warrant an evacuation that evacuees start evacuating. It is a general misconception that people panic and stop acting rationally as soon as they see a situation like a fire. Several studies [17–22] have disproved this claim. Irrational panic is hardly ever the standard first reaction to seeing fire. Rather, people react rationally and try to gain more information so that they can act in a manner more appropriate to the situation.

This section presented the literature on how the process of evacuation starts. The process called *pre-evacuation* consists of two phases: perception and interpretation. Cues present in the environment indicate to the evacuees that there is something happening in the environment. These cues can be of different types and their effect varies based on certain factors of the environment and certain characteristics intrinsic to the evacuee. During the second phase people search for verification that the situation does indeed require some action and various cues and other factors help him towards this realization.

2.2.2 Crowd behavior

This section discusses different theories of crowd behavior that, in general, deal with the phase that commences after *pre-evacuation*. According to Tong and Canter [47] emergency egress is a *complex social process*. What this complex social process is, has generally been a matter of much debate with a great deal of conflicting theories on the same. In his PhD thesis [20], Torres categorizes the theories on crowd behavior into six:

- **Social Breakdown Model:** According to this theory, a fire evacuation is characterized by competitive behavior (pushing and shoving). As competition develops traffic jams occur at the doors or passageways. This competitive behavior

is because the individuals think only about getting themselves out without concern for the other trying to escape, low optimism and a large group size which makes people think that they need to push and shove to stay alive. This model suggests that the reason for deaths and problems during egress is because of *non-adaptive behavior*. Non-adaptive crowd behavior is the type of crowd behavior that does not adapt to an emergency situation and often leads to destructive consequences.

- **Hysterical Belief Model** : This model is also characterized by flight and non-adaptive behavior. However its cause is slightly different. In this model, it is necessary for a collective belief in a generalized threat to develop before people start panicking. This belief in a generalized threat, develops when people perceive ambiguous cues and then feel anxious due to certain factors, like a feeling of reduction in exit choices; this, in turn, leads to a feeling of unambiguity in believing in a generalized threat which they can counter by fleeing. Also certain *components of social action* need to have been completed that permit panic to occur. These components are characterized as being actions that are as per norms (i.e. not non-adaptive) and which are done for the benefit of the entire crowd or collective. Emergent norm theory can be considered to be a kind of hysterical belief model. According to this, there is an extended period of *milling* during which people engage with others and exchange experiences. During this period a consensus develops and leaders also develop. Everyone in the group follows the actions of the leader in a process referred to as *keynoting*.
- **Non-Social Model** : This theory's main proponent is Quarantelli [54]. According to this theory, people largely act in a pro-social manner; yet there occurs a point where a collective threat becomes severe and conditions emerge where social bonds break down and actions designed for self preservation emerge. The key point of difference from the previous models is the importance given to social bonds and the association of flight or competitive behavior or panic with breaking of these social bonds (this is a necessary but not a sufficient condition). Also competitive behavior is associated with individuals and not the group as a whole as specified in the hysterical belief model. As above, a perception of dwindling number of exits is also a precondition.
- **Normative Model** : This model is characterized by people helping one another during the emergency egress and an overall lack of competitive behavior. This theory is also characterized by a tendency to fight fire.
- **Affiliative Model** : This model conceptualizes emergency egress as people engaged in stimulation seeking behavior by moving towards others they feel close to or running towards locations that are familiar. People are assumed to have cognitive maps and a cue triggers a tendency to try to restore congruity to

their cognitive map. People seek familiarity to such an extent that they might even move towards a threat. Even competitive behavior is affiliative in nature, i.e., people try to exit with their primary groups rather than alone (as in the non-social model). This search for familiarity might manifest itself as a search for group members or in other cases as taking a familiar route or in others just moving towards locations that are familiar.

- **Self Categorization Theory** : In this model, groups formed aren't fixed. People tend to develop a sense of 'we-ness' with others present in the evacuation and form groups dynamically. Drury et al. [55] showed that collective bonds may be strengthened and even created through the experience of an emergency. This is characterized by people helping strangers escape. This dynamic emergence of groups was also observed by several others [19, 33, 53, 56]. When part of a group, people tend to follow the group consensus. Leaders develop for groups and these leaders can impose their will on the group. Groups that are formed are highly dependent on context and other groups and in turn the individuals present. Groups aren't static, they have dynamic characteristics. When faced with an adversity (like a fire), behavior in a group reaches a common level with time. The elaborated social identity model (ESIM) for crowds was an extension to this theory proposed by Reicher [19] which proposed an even more dynamic model of social identity and group behavior where the social role or identity of an individual can be influenced by his actions and vice versa.

Tong and Canter [47] stated that there are three distinct strategies that people adopt in a fire, the first of which is to (try to) extinguish the fire, the second is to seek shelter and wait to be rescued and the third is evacuation. Most of the theories state that people engage in flight behavior, i.e. they attempt to escape rather than fight the fire. Only the normative model predicts a fire fighting behavior and this kind of behavior was shown by only 4 of the people surveyed by Torres. It is interesting to note that in over three quarters of the domestic fire incidences that occurred in the United Kingdom and Australia, the fire was extinguished either naturally or by the residents; fire fighters were not called [17]. However, as further highlighted in the same paper by Kobes et al. [17] there is very little literature on fire fighting behavior shown by occupants besides this. This might be because of a difference in attitudes towards private and public property. Nevertheless, there is not enough evidence to show that fire-fighting is normally done by occupants. This, however, does not suggest an absence of altruistic behavior. Other models like the affiliative model and SCT model do suggest that people do show altruistic behavior in helping others who are part of their group or who they characterize as being part of their group.

All the models predict that as long as people are capable of escape, they try to escape; and panic is not an influential factor. The chart shown in Fig. 2.1 (obtained from [13]), shows the relationship between mobility, complexity of the architec-

Figure 2.1: According to this chart from [13] there are four strategies that humans resort to in egress. Studies have shown though that more often than not, egress or slow egress are the only strategies taken.

ture and people's reaction to fire. This chart indicates how, unless forced to, people do not defend in place or try to find refuge. It also confirms the preference for familiarity (of a refuge) to staying in place and defending against the fire. The fact that people always try to escape and rarely seek shelter even in the case of extreme smoke has been confirmed by Kobes et al. [17].

However, it is important to note that flight behavior is not equivalent to competitive behavior characterized by pushing and shoving. The competitive behavior, which is central to the first two models listed above, has been established to be non-existent during egress, at least not in the sense of it spreading through the crowd. Cocking and Drury [42], who are the main proponents of the self categorization theory, do admit that individual panic does occur; but rather than the panic spreading through the crowd it is the calm people who manage to impose their will on other people and calm the people who panic. The same thoughts were echoed in other studies [18, 21, 22]. Several studies [21, 22, 55] actually take a much stronger stand that competitive behavior almost never happens and rather people almost always behave altruistically. This was further confirmed by Torres [20] who found that individual competitive behavior never occurred. Nevertheless, people did do whatever they could to enhance their and their group's survival but without knowingly causing harm to others (which is what competitive behavior/ panic entails).

A further point in favor of the latter four models is the importance that they give to social bonds. For example [36, 57, 58], among others, have already demonstrated the important role that familiarity plays in behavior during egress. This was reaffirmed by Torres [20], who found that even in cases of extreme danger, people still maintained their social bonds. In fact, as opposed to the non-social model, in the real life evacuation, social bonds never broke down. If people did manage to get out of the fire, they would try to come back or help their trapped group members in some way [17, 20].

The normative, affiliative and SCT models also recognize the importance of social roles in determining the reactions of people to fire. Several studies have made this observation [21, 37, 43, 46, 47, 53]. These models explain how the employees of a workplace who are trained or prepared, generally try to guide customers to the closest exits. They also indicate that within groups, the leaders or caregivers (e.g. teachers or parents leading a children) continue to take the role of leaders during egress and try to find safe evacuation paths to protect their group.

In summary, most experiments and studies have been in favor of an affiliative or SCT based model with certain elements of the normative model included. The consensus is that people always try to escape while trying not to compete with each other and they look for familiar people or locations in dangerous situations and if possible they even form new bonds with strangers stuck in similar situations. Also, the social role of a person determines his behavior. Staff try to ensure safe egress of customers and social leaders (parents, teachers, etc.) take more responsibility during these times. Irrationality of occupants and panicking crowds are some of the common themes running through news reports of fire incidents. This is discussed in more detail in the next section.

2.2.3 On the rationality of crowds

It has been shown [17–22] that very rarely do people panic and behave irrationally in spite of time constraints and stress. However, situations have been observed where people do not take the best possible actions and do not act in their best interests [37]. This is probably because a person’s rational thinking is bounded by the constraints of his/her knowledge of the environment i.e. he/she is only bounded rational. Bounded rationality is an idea that is becoming increasingly popular in behavioral economics and it implies three major things [59]:

1. People do not know everything.
2. People cannot think infinitely into the future. They think in the short term and tend to take decisions that benefit them in the short term.
3. People are generally loss averse i.e. they do not like to take a risk of loosing more even if the gains are much higher. In a fire evacuation this means that

people avoid action for small problems and take drastic actions as things get out of control. This is because early action might result in huge losses [48].

If they are only bounded rational and they tend to make worse decisions as the fire gets worse, then it implies that the time pressure and stress must play some part. Ozel [26] proposed a decision making theory which explains this. The basic premise of his theory is that given the same set of information, people may attend to information differently depending on the stress and amount of time pressure they experience. When people gain information they feel less stressed. This can also explain why people investigate and try to gain more information when they perceive the first few ambiguous cues of fire.

One of the significant ideas suggested by Ozel [26] is the idea of filtration. According to this, to cope with time pressure and unavoidable conditions, people try to filter out all the cues that they deem as unimportant or less important. The rate of this filtration increases as the time pressure increases and as a result cue utilization decreases. This in turn implies that they have lesser information about the environment and tend to behave less rationally. This effect has also been observed by Davis et al. [27] in older people who, due to age, are only able to process less information. The significance of information in egress and more particularly perception will be discussed in more detail in Chapter 4.

Thus it is understood that humans act rationally within the bounds of their knowledge. However, as the situation worsens, their behavior tends to look more irrational to an observer. This is because the evacuees observe less on being stressed and as a result have lesser knowledge; this subsequently results in, what appears to be, irrational actions.

2.2.4 Summary

The previous pages of this thesis have discussed the existing understanding on human behavior during egress through the analysis of incident reports, experimental studies and other reports in a wide number of fields. Several features of human behavior during egress that could play a key role in ensuring safe egress were discovered through this analysis. These are briefly summarized below:

- Ambiguous cues: Most cues are ambiguous in nature and they need to be present for a certain period of time before catching someone's attention enough for them to even start investigating.
- Early investigative behavior: Early movement is characterized by investigation rather than egress. The investigators then gain knowledge or information either from what they observe directly or from what they hear from others. So there is a *spread of information* in the environment [50, 60].

- Flight behavior: Once people interpret the situation as fire, they react to it by trying to get out of the building. They do not try to fight the fire or to seek shelter unless they are forced to by the circumstances.
- Search for familiarity: People's first action while trying escape is to move around and try to find their primary group members who they are familiar with. They then try to exit through some exit that they are familiar with rather than observing new signs. The search for familiarity, also implies the importance of groups and social bonds, to the extent that even after exiting the building, people might come back to help people from their group.
- Absence of competitive behavior: People do not engage in competitively pushing or shoving others when engaging in egress. The rare cases when pushing or shoving occurs, it is because the individual involved is forced to do so.
- Altruistic behavior is common: People tend to help others who need help. They also do not panic or act irrationally. Instead, they follow social norms and try to be as orderly as possible.
- Decisions and behaviors are dynamic: As people interact with the environment and the other people in the crowd their decisions and behavior might change or be influenced.
- Bounded rationality: Stress and time constraints cause a reduction in cue utilization and information availability and thus cause people to make decisions that with complete information will seem irrational.

The next section presents an analysis of existing approaches to computation modelling of egress. These models are analyzed against the understanding of human behavior that has been presented in this section. This helps in identifying the research problems that are tackled in the rest of the thesis.

2.3 Computational Models of Egress

Conducting experiments and analyzing an environment or an event for its safeness can be expensive and, more importantly, quite dangerous. Through modelling and simulation, it might be possible to prevent the unnecessary loss of life. It can help architects design buildings better and the management and fire fighters to handle situations better. These models have to be able to accurately simulate how humans behave in order to be useful. A good predictive model can also help test the existing theories of human behavior through the study of emergent behavior. In the previous section the reader was introduced to psychology and sociology literature on how humans and crowds behave in fire and other evacuation scenarios. In this section,

computational models of egress that have been developed over the years will be discussed. Section 2.3.1 introduces the different components that generally constitute a computation model of egress. Section 2.3.2 presents a broad overview of the different kinds of modelling approaches along with some of their strengths and weaknesses. Finally, Section 2.3.3 explains in more detail some selected models.

2.3.1 Components of a computational model of egress

Depending on the purpose of the model, the approach used and the level of detail, a model can consist of various components and sub-components. At the most basic level each model will have: a model of the environment and a model of individuals involved.

2.3.1.1 A model of the environment

Most models have an environmental model that represents the physical environment of the location where egress is taking place. They may be continuous or discrete and generally consider a single floor with passageways separated by walls. Doors are almost never modeled. At a lower level, depending on the scale of the simulation, in-room obstacles like pillars or furniture are also sometimes simulated. Sometimes, models do consider multi-storey environments. However, movement along staircases and escalators is more complicated and generally ignored except in rare cases [32, 61].

An important point of difference is generally the resolution of the environment. Coarse network models have entire rooms as their smallest entity and cannot predict within room complexities of movement, while others have fine networks, which can do this through the division of the rooms themselves into fine grids. Even this minor approximation is avoided in models that use continuous space. These models can thus model the dynamics of the interaction between evacuees.

It is also useful to have some model of how the fire or smoke will spread within the environment. In some cases, this can be done by creating a separate external model of the fire or smoke and importing it for the simulation or it can be done in the same model and simulated in parallel with the model of egress. Olenick and Carpenter's report [62] gives a compilation of different models that are used for modelling and simulating fires. In general, however, calculating smoke spread in a simulation is computationally expensive and is often ignored.

2.3.1.2 A model of the individuals engaging in egress

The model of the individual is the most important part of an agent based crowd model and there are a wide variety of ways in which this can be done. These will be discussed in more detail in Section 2.3.2 and Section 2.3.3. There are a variety of things to be considered when a human is being modeled. These include:

- **Physical Representation:** This refers to the physical characteristics of the humans being like the shape and size of the model used. This is discussed in much detail in [3, 63]. Some papers suggest that for accurate modelling an elliptical shape is best but to make this computationally efficient a 3 circle model can also be used like in [63, 64]. The speed of movement of the humans and the time taken for pre-evacuation behavior can also be considered to be part of the physical representation. This is also extensively discussed in some surveys [49, 60].
- **Navigation:** Navigation refers to how the agents move within an environment. Depending on the scale of the environment, this generally consists of a higher level path planning which is generally A-Star and a lower level collision avoidance algorithm. The choice of collision avoidance algorithms can have significant effects on the dynamics produced and this is discussed in more detail in Chapter 7. Movement during egress can have complications because of the effects of fire and smoke on visibility and following behavior [17, 65, 66]. The partial knowledge of evacuees can have an impact on higher level path planning.
- **Knowledge:** This can refer to either knowledge of events or knowledge of the environment/ layout. Most models assume complete knowledge of both. However, there are some that do model individual specific knowledge and exchange of information. These will be discussed in more detail later sections.
- **Behavior and decision making capacity:** This refers to the detail in which some of the behavior mentioned in Section 2.2.4 is modeled. This also refers to the social interactions that takes place between the evacuees. Some models do not consider behavior at all, while others have sophisticated models of decision and behavior. The same behavior can sometimes be produced by using different techniques. Some models use a functional analogy, like social forces, which approximate behavior through mathematical formulas; while others use rule based techniques or more complicated hierarchical models. These will be discussed in the following sections.
- **A model of trained staff:** It is generally recognized that the employees of a place or the people trained in handling a fire evacuation have a key role to play during egress [21, 23, 36, 43]. Some models do take this into account.

In the following sections, some of the different approaches that are used for modelling crowd evacuation are presented. Not all models have all the components mentioned above, but most models have at least a few of them.

2.3.2 An overview of existing computational models

There are several reviews of computational models of egress. Some [5, 18, 67–69] list and attempt to categorize existing models according to their properties. The literature review in Still’s thesis [3] and Santos and Aguirre’s article [23] give comparisons and critical analyses of some of the most popular computation models of crowd egress. The former gives an analysis from a computational perspective while the latter analyses the models from a psychological perspective. Together they form an invaluable source of information on computational models and their strengths and shortcomings.

The different behavioral criteria along which computation models can be differentiated are discussed in some surveys [5, 68]. Others [18, 69] give a comparison of the models based on an engineer’s or a modeler’s perspective. It gives an idea of the computational complexity of each of these models. Schadschneider et al. [18] explained strengths and weaknesses of each approach in more detail, while Zheng et al. [69] lists the different models of each kind and their features.

As mentioned in Section 2.1, all these differences can broadly be said to be differences in level of detail. At the most abstract level, we have models that have a coarse network of just rooms connected to other rooms and homogeneous uniform crowd behaving like a fluid. At the most detailed level, we could have an agent based model that considers each individual to be different from every other and with each having its own memory, decision making capacity and behavior, all of which are in turn influenced by their interaction with other people in the group.

In the following pages, some of the typical models of each approach are introduced and analyzed with respect to the discussion in Section 2.2.4.

2.3.2.1 Network or Queuing Theory based approaches

Network or Queuing Theory based approaches use the coarse network approach mentioned earlier. Nodes are used to represent rooms and passageways and generally any place which can hold people. The arcs in the network represent the gateways or connections through which people move from one node to the other.

One of the earliest models of this kind is Evacnet+ [70]. In this model, the user specifies the capacity and initial content for each node and the traversal time and arc flow capacity of each arc. Queuing theory is then used to calculate the time required to evacuate the building. As in other networks, waiting time, throughput, length and utilization of each arc can also be determined. [71] is another example of a similar queuing theory based approach which has a basic ability to show heterogeneity. But this ability is very limited because there is no way that grouping behavior, investigative behavior or effect of time and stress can be modeled. More importantly they do not correspond to a tangible reality [72]. [73] is another example of this approach.

2.3.2.2 Flow models

In flow models the crowd is assumed to be a fluid and the crowd motion is predicted based on the geographical layout, density and velocity of the particles. Schadschneider et al. [18] provided an excellent analysis of these kinds of models. The first of these models was Henderson’s model [74] according to which the interactions between the pedestrians were calculated using the kinetic theory of gases. One of the theory’s main drawbacks was the assumption of energy conservation and Newton’s third law being applicable to crowds.

These drawbacks were removed in later models which made use of a density function. This function was derived from Boltzmann’s transport equation that describes the change for a given state as the difference of inflow and outflow due to binary collisions. These later models were also able to distinguish between different groups of particles that had different destinations. These models however fail to work at low densities and cannot model the complex heterogeneity of a crowd [72].

2.3.2.3 Environment Control Based Models

These models are more complex than the previous two models and generally manage to model some complexity of behavior. The idea behind these models [75] is that it will be computationally difficult to model a very large environment with a large number of people, each of them instilled with some complex decision making and behavioral ability. So rather than model complex entities, the decision making ability and “behavior” of the agents are stored in the environment itself. The location of an agent determines its behavior. Thus, these models can be termed to as being controlled by the environment.

Banerjee et al. [75, 76] propose an environment with several different layers that together store all the information that is needed for the simulation. This reduces the complexity of the agents considerably. The idea of using environments for computation with the agents having limited complexity is one of the most established in the field of modelling and simulation. Cellular Automaton models and lattice gas models have been around for a long time, and a majority of the models of these kinds use what can be termed as *environment control* to implement complex behavior.

2.3.2.3.1 Cellular Automaton based approaches: A Cellular Automaton (CA) model is one in which space and time are discrete. The state space of a CA model is also discrete and finite. In each time step, the values of all cells are updated synchronously based on the values of cells in their neighborhoods. Depending on the type of neighborhood (i.e., von Neumann, Moore), and the type of lattice (triangular, square, hexagonal, etc.), the exact number of cells in the neighborhood of a given cell can vary [77].

This technique is very popular and has been used in a number of different ways. In [5, 68] these models are called *fine network models*. [78] is a typical, simple CA model in which the agent position is updated based on its distance to each possible exit and the density of the crowd at each possible exit. The preference for familiar exits and grouping behavior is also modeled. However, it was used only for modelling single room evacuation.

One of the drawbacks of a typical CA model like this is that movement is simple in that agents can only move to one of the fixed neighboring cells which are present at fixed angles. This is limiting when trying to model detailed motion. In [79] a real coded cellular automata model is proposed. It is called a real coded cellular automata because of its ability to consider velocities at angles other than multiples of 45 and magnitudes more than a single cell length away (non discrete velocities). Each time an agent's location ends up at a point other than one of the grid points, it is assigned one of the 4 neighboring points probabilistically. Some approaches [80] use a hexagonal grid layout instead of square cells. This gives more freedom of motion.

In [32] slightly more heterogeneity can be modeled since each agent can have a different velocity. Though it is just a movement model, the pre-evacuation time is modeled through the use of a delay parameter specified for each agent. The same thesis suggests that by combining a CA model with a network based model, these models can be used for modelling egress from complicated buildings also.

Most simple CA based models tend to concentrate on movement and avoid modelling complex behavior.

2.3.2.3.2 Lattice gas models: Lattice gas models are CA models that make use of a discretized version of the Boltzmann transport equation to model motion [66, 81, 82]. These models do not generally model any complex behavior. In fact they are generally used on a very small scale to study patterns of egress from a single room and to study the effect of different obstacles, number of exits, etc. on egress in extremely dense environments.

2.3.2.3.3 Floor field models: Floor field models are more complicated CA models that are more capable because agents can send messages between each other and communicate.

The Extended Floor Field Model [83] in which agents interact through virtual traces that act like the pheromones in chemotaxis and the SWARM information model [84] which uses multiple floor fields to model transmission of knowledge between agents are typical examples of floor field models. In the SWARM information model, there are multiple static fields on the floor to indicate the different world views. A higher number indicates a more accurate model. At certain points (for example, points where maps are displayed), the agent can upgrade its knowledge to a more accurate model. It also models communication at a basic level. When an agent with

a higher index comes in contact with one with lower index, then the lower agent's index is upgraded. Thus it has a basic method of communication and exchange of information. [85] is also a similar model that has a static potential field guiding towards the goal and a dynamically generated interaction field.

The Situated Cellular Automata (SCA) modelling framework [86] was proposed so that floor field models can be extended to model more complicated behavior by psychologists and sociologists based on their findings. The authors call it a multi-agent systems based approach, but most of the calculation and processing is done by the environment and not the agents, hence we call this also an environment control based model. Each agent in this model has a particular state that it is in, which influences the way it acts. Agents are allowed to emit and store messages in specific locations. When agents reach a particular location they react to all the messages present there (some from other agents, others from the environment). In this way communication and group behavior can be modeled effectively.

While similar models (especially the SCA model) might, in the future, be extended to model complicated high level behavior of the kind mentioned in Section 2.2.4 it is not a very intuitive approach. It is more natural for people to think of behavior at an individual level. Programming behavior through an environment requires abstractions that are difficult for all but the most experienced modelers.

2.3.2.4 Agent based approaches

Agent Based modelling is the preferred approach for modelling complex behavior [87–89] because the bottom up approach makes it easier to view and break down problems into more manageable entities. Woolridge [90] defines an agent as a computer system that is situated in some environment, and that is capable of autonomous action in this environment in order to meet its design objectives. He further defines *intelligent agents* as those agents that are capable of flexible autonomous action through reactivity, proactivity and social ability. The agent architecture is the software architecture used for modelling this decision making ability. Specifying an architecture would provide a structure that can be used to break down the complicated process of decision making. There are several frameworks like the Belief-Desire-Intention (BDI) framework [91] and the Recognition Primed Decision (RPD) that provide a structure on which agent based architectures can be developed.

An example is [92] which is a NetLogo based agent based simulation which modeled the effect of *management* and trained staff. This is done by programming officer agents that instruct other agents to leave and give them instructions on how to do the same. This is a relatively simple task in agent based modelling; models that simulate more complicated emotions, behavior, social interaction and decision making are discussed in more detail in Section 2.3.3.

2.3.2.5 Hybrid approaches

Most models use one of these different approaches or can be classified into one of these. However, there are certain exceptions that can't easily be classified as any of these. For example, the pedestrian motion model proposed in [72] is a multi agent model that uses a cellular automata at the lowest level for collision avoidance. In fact, it is not even a simple CA model, because it uses a network based representation for path finding and a radial individual specific discretization of space to capture decisions about the direction of walking and the aforementioned CA layer at the bottommost layer.

2.3.3 Detailed discussion of significant computational models

The previous section introduced some of the popular approaches to crowd simulation. In this section, models that manage to simulate some of the key characteristics explained in Section 2.2.4 are introduced.

2.3.3.1 Pires's model of pre-evacuation behavior

Unlike the other models that are described in this section, Pires's model [31] is not a complete model of egress behavior. Nevertheless, it is important because it proposes a method by which the pre-evacuation decision making of an individual can be simulated using a simple Bayesian Belief Network (BBN). It also takes into the consideration the effect of time constraints and stress on the BBN.

However, this approach has its limitations. Most of the details of the pre-evacuation behavior are abstracted away as probabilities which are estimated by *experts*. So, besides not being able to simulate the effect of different kinds of cues, this model does not make it easy to analyze and understand the behavior and movement produced.

2.3.3.2 Legion

Still's Legion model [3] is based on extensive analysis of crowds exiting stadiums in the UK. The Legion model was extended from the Vegas Model, which worked on the basis of an extensive set of rules that governed the behavior of each agent. Such an approach, Still found, had two basic flaws:

- One cannot determine, a priori, all the possible conditions that can possibly occur. The same person might have different reactions on different days. As a result, a lot of these specific rules and conditions could be replaced by noise.
- A lot of behavior can actually emerge due to the self organization of the system, without actually specifying these rules.

Still realized that it would be impractical to consider a parameter for smoke, another for nature of threat, another for the emergency and so forth. According to him, during an emergency, a human either moves towards the threat (investigate), stays in place (ignore) or moves away from the threat (evacuate). This choice is based on the value of three parameters that interact- Objective (try to move to desired or intended end point), Motility (try to maintain optimal velocity) and Constraint (try to maintain a minimum distance between yourself and the other objects)- and one parameter that represents the reaction time- Assimilation (delay in reading and reacting to the environment). He has explained how all the key factors in an evacuation, for example, communication, alertness, social role, position, effect of location, population density and so forth, can all be modeled using just these four factors.

The model is notable for presenting various studies to accurately model the size and shape of human beings. However, it comes to the conclusion that a square cell will most accurately model a human being considering the possible ways in which he can turn. It is interesting to note that others [63, 64] however, have shown that an elliptic or a tri-circle model of humans would be more accurate.

Still uses environment control, through what he calls iSpace, for modelling communication between agents. However, the perception model used is quite simple and the idea of knowledge spread and partial knowledge is not considered. As a result, even though the idea of emergent behavior is captured brilliantly in this model and a lot of emergent natural phenomena are obtained, it has its limitations like the inability to model pre-evacuation behavior and partial knowledge.

The greatest strength of this model lies in the importance given to data, both for theorizing and validation. Also, the physical model used, which was based on an extensive study of the existing data, is something to be emulated. Most importantly, the idea of emergent behavior in crowds was emphasized by the model and experiments. His adherence to Occam's razor in keeping the model as simple as possible and letting more complicated behavior emerge from this is a further strength of this model.

2.3.3.3 MACES + PMFServ

This model [38] is most notable for the approach taken to modelling complex behavior. In order to obtain believable emergent behavior a popular psychological model of human beings (PMFServ) was integrated into their existing Crowd Simulation System (MACES).

The original MACES model was a multi-agent system based model, with each agent having its own representation of the map. Through exploration and communication they create a more detailed view of the world. This, combined with a model of leadership and non leadership behavior, formed the high level decision making which gave a destination to the lower level path planning algorithm. This path

planner then used Helbing's social force model [10] for collision avoidance. This is one of the few models which takes into consideration the partial knowledge of occupants. However, it is assumed that once informed about a room, the agent memorizes and never forgets it.

PMFserv was conceived as a software system that would expose a large library of well-established and data-grounded Performance Moderator Functions (PMFs) and Human Behavior Representations for use by cognitive architectures deployed in a variety of simulation environments. Its principal feature is a model of decision-making based on emotional subjective utility constrained by stress and physiology. The basis for this decision making model is Maslov's hierarchy of needs [93] and the OCC model [94]. Needs reservoirs corresponding to the degree to which the agent has satisfied his needs are set based on any action that might have occurred in between decision cycles. The OCC model describes a hierarchy that classifies 22 emotion types. Using this hierarchy, the mood of a person in a particular situation can be determined. Besides PMFServ, there are other approaches to emotion modelling for agents as well. For example, appraisal theory based approaches are also used in certain Agent Based Models [39].

In summary, in this architecture, PMFServ provides the emotional and behavioral basis which is utilized by MACES for motion. Though it has a model of behavior, emotions and decision making it does not model pre-evacuation behavior or the search for familiarity. However, it is still of interest to us because of the model of partial knowledge and the idea of having a dedicated and extensible module for behavior with an extensive database of emotions and their effects on decision making. This sort of extensibility is very useful for a behavior model of crowds. In fact, a crowd simulation framework should be flexible enough that different theories of behavior can be used without too much effort. However, one drawback of the architecture used here is that emotions and behavior have no effect on the lower level social forces model. This is a natural difficulty of this approach of improving the model by integrating another existing model instead of having a single central model.

2.3.3.4 Collective Panic Behavior Model

França et al. [41], created a simulation model of panic behavior which was based on the hysterical belief theory explained in Section 2.2.2. To recap, this theory believes that people normally follow social norms and behaviors. A fire alarm or smoke causes social unrest. Following this, people start investigating and communicating with each other and start forming a consensus about what this threat represents. This process of interaction is called *milling*. Following this, a stage of collective excitement is reached wherein a *collective belief (CB)* is formed between the people and the people no longer have an individual thinking capacity. Next, a social contagion stage is reached wherein the CB is highly contagious and spreads to other people who come in contact with

this crowd. After a short period of time, this situation escalates into collective panic where none of the actions of the crowd can be accurately predicted because they are believed to behave irrationally.

Most aspects of this theory of panic have been proven to be wrong as explained in Section 2.2. Nevertheless, this model is still interesting due to the detailed way in which the behavior model has been converted into a computation model. Also, the agent architecture is notable for the possibilities it holds for extension.

Each agent is made up of four modules: a Belief and Knowledge Management Module (BKMM), a Perturbation Module (PM), Dissipation Module (DM) and Social Cognitive Module (SCM). The environment itself is divided into three parts: a physical environment, a communication environment and a group mind representation.

- The PM is responsible for picking up cues from the communication environment and analyzing them. These cues are first semantically analyzed using rules specified in the BKMM to determine the mood of the message (loving, aggressive, neutral). Once analyzed, the information is stored in the BKMM as either a belief or as knowledge based on the amount of evidence for this belief.
- On obtaining evidence, beliefs may be changed to knowledge. The knowledge base has three parts to it: The first represents the intrinsic features of the agent (e.g. health, speed), the second represents the knowledge that it gains through perception (e.g. temperature) and the third represents social state variables that are determined by its interaction with other agents (susceptibility to other agents). The BKMM also has a rule base that stores information on how to behave and act and interpret information. The rule base itself has three parts: a functional part which gives the agent an identity and goals and instructions to pursue them. A Dynamics module which relies on learning and a reactive module related to the agent's survival which is usually time-constrained. There is a micro representation of the CB in each agent. According to the crowd theory used, there is a tendency for this belief to move towards the macro CB.
- Social Cognitive Module is like the control unit of each agent. This module again consists of three parts: Firstly, it has a Cognitive Core module (CogC) that continuously processes and manages information and guides actions so that the agent can pursue and fulfill its goals. An event that poses a threat triggers the Collective behavior core to take over. This makes the agent act in a collective manner. The trigger is done based on experiences stored in the memory of the agent. Uncertainty about information, due to only partial knowledge, triggers social unrest and milling. During this process, the agent builds up its micro CB representation. As the panic rises, the dynamic rules become less important. The knowledge base has a variable for permissiveness, which represents the

adequacy of the actions of the group to the situation. When the actions cross a threshold, i.e. they are inadequate for the situation, collective excitement results. At this point, even the functional rules start playing and social norms start being broken. After a certain point, this situation escalates into panic and all the agents act according to the macro CB which enforces a competitive behavior among the agents. The third part of the SCM called communication core is responsible for dispatching instructions from the CogC or CBC to the dissipation module through a queue.

- The DM is responsible for sending messages into the environment in the correct format according to the Agent Communication Language. It does so on getting instructions from the SCM and adds details to the message like the mood of the message and whether the information is conveyed as a gesture or a speech.

Detailed communication between agents is possible through the use of ACL with details of mood and the type of gesture. Agents don't directly exchange information. Instead, the communicating agent dissipates the message into the environment; and autonomous agents check the communication environment in order to check whether the message is relevant for reaching its goal. The BKMM module represents the idea of information processing and uncertainty reduction accurately. However, the theory of panic and collective behavior, on which the model is based, is much debated [20, 24].

Combined with the PM, a model similar to this can in theory be used for modelling all the pre-evacuation behavior that is necessary. While it is possible to make all these extensions, it doesn't seem like the authors have attempted to do this in their model. As expected, there is no affiliation behavior as the Hysterical Belief approach on which the model is based has no concept of social bonds and their importance. The absence of altruistic behavior and the importance given to competitive behavior are some other drawbacks. The major drawback, however, is the lack of details about the physical model of the agents or the environment or the level of detail used.

2.3.3.5 Exodus

Exodus [95] is one of the most detailed and mature evacuation simulation tools available. It has different packages for building, maritime and aircraft environments. The buildingExodus model comprises five core interacting sub-models: these are the occupant, movement, behavior, toxicity and hazard sub models (see Fig. 2.2).

Once the building layout is specified as either a DXF file, CAD package, or by using the interactive tools provided in the suite, it is converted into a spatial grid representation which stores all the details of the building including obstacles present inside it. It is also able to store and show movement at staircases and across multiple

Figure 2.2: Interaction of EXODUS modules [96].

levels. The grid itself is represented as a graph with the nodes representing a small region of space and each arc representing the distance between each node. Individuals travel from node to node along the arcs.

Each individual has more than 20 behavioral attributes that are specified at the start of the simulation and stored in the occupant submodule. Based on these, the behavioral module determines the behavior of the agents. Two kinds of behavior are modeled: *normal* behavior and *extreme* behavior. The only difference between the two is that in extreme behavior, the agents have a patience threshold beyond which they will stop queuing and resort to some other course of action. The behavior model works on two layers: a global layer and a local layer. The global layer gives a set of waypoints or a single destination (possibly the exit or a familiar location). The local behavior is determined by the agent's attributes and determines factors like how long he waits before evacuating, conflict resolution and other things. Some of these factors and local behavior are probabilistically determined.

The toxicity submodule determines the physiological impact of the environment on the occupant. This submodule assumes that the effect of smoke or fire is determined by the dose received and not the exposure concentration. The toxicity affects the occupant's mobility agility and travel rates. The hazard module on the other hand is responsible for generating hazards like smoke and fire in the environment as a function of time and location. A dedicated zonal fire simulation model called CFAST is used for this purpose. Several analytics tools have also been developed to analyze the results produced by the simulations.

Gwynne et al. [96] added some new features including the ability to

specify occupant specific knowledge of the layout of the building. The behavior module is further modified so that the agent chooses to use emergency exits that he knows of only when he is showing extreme behavior. This representation of a map of the place rather than just a path to exit allows the occupant to dynamically re-plan his route. He is also given the ability to learn from signs and through interactions with other agents.

The effect of smoke is also modeled by reducing the efficiency of the path selected and reducing the speed of movement as per the concentration of the smoke. This reduction in efficiency and speed happens only as long as the exit is not visible to the evacuee. A certain inertia in route replanning is also modeled to indicate the reluctance that people have in changing their pre-decided route.

Hollmann et al. [97] presented a theoretical prototype of an emotional model that modeled the effect of time pressure and stress on evacuees was introduced into buildingEXODUS. Each agent is specified by giving a list of tasks or waypoints that he has to visit before or during the evacuation. The tasks are categorized as being compulsory, time critical and elective and each of them are given an estimated time for completion. An urgency factor indicates how constrained for time the agent perceives itself to be. The speed, drive, patience and the itinerary of each agent is rearranged based on the time elapsed, time left and the feelings of the person. Without an actual implementation of the prototype, it is difficult to analyze the strengths and weaknesses of the proposed emotional model.

The strength of EXODUS is the fact that it has a dedicated fire/ smoke simulation engine and a toxicity calculation module. The prototype of the emotion model to be used is also an innovative addition. The ability to specify so many parameters and thus the heterogeneity of agents is also important, as is the ability to model complex building layouts and obstacles in a variety of ways. Nevertheless, the model falls short on a few key factors. It does not try to model crowd behavior or its effects. The idea of cue perception and pre-evacuation behavior is very limited, unless the user specifies a list of tasks that each occupant is to do [95]. Moreover, the theoretical prototype for doing this was never implemented and tested.

2.3.3.6 ESCAPES

ESCAPES [29] is a multi agent based evacuation simulation software tailor-made for airport environments. Unlike most of the other models mentioned earlier, it also has a 3D visualization engine, which makes it more usable for security purposes for non experts. One of the ways in which this model improves on others is that it can model families and social bonds. It models how evacuees search for their family members before attempting to evacuate. Family members also communicate all their knowledge to other family members. It also models a fear factor and emotional contagion. The knowledge of the fire spreads and the authority figures are able to calm people down.

The model also takes into consideration the incomplete knowledge of most people and the spread of knowledge in the environment. However, the models for these are simplistic with only 2-3 states of fear, emotion and knowledge. The reaction time is modeled as a function of the proximity to the event. Agents that are near an event evacuate immediately, whereas others, will behave normally until they get enough information to know that they need to exit.

While the model includes some behavior like affiliation (within families), the prevalence of competitive behavior in the model is unrealistic. The family sizes are fixed (parents and two children) and the only social bonds that exist are within this family. The model of partial knowledge is rather simplistic as is the investigative behavior model. Despite these drawbacks, this is one of the few models that consider almost all the factors discussed in Section 2.2.4.

2.3.3.7 MASSEgress

Pan's [40] multi agent based model of evacuation called MASSEgress is notable for the well structured and detailed architecture of the agents used and the importance that he gives to non-adaptive behavior being synonymous with a high stress evacuation environment.

Mobility, age, gender and body dimension are the only intrinsic characters that are included. A population generator is used to generate the required distribution of people based on these factors. The simulation uses a grid based physical environment which can be loaded as a CAD file. The data from simulation is logged to facilitate easy analysis, and a 2D and 3D visualization engine is also implemented for ease of use and simple analysis. The simulation engine itself simulates the behavior by dividing it into three parts: a perception engine, a behavior engine and a motor engine.

The perception model used is relatively complicated compared to the models discussed earlier. It consists of a point test method and a ray tracing algorithm implemented along with a visual cone. The visual cone specifies the region in which the agent can perceive and is determined by the eye position, viewing angle and visual range of the agents. The point test algorithm determines whether an exit or waypoint is perceived. The ray tracing algorithm determines which of the obstacles are perceived and ensures that only objects closest to the agent are perceived. modelling only the visual component of perception limits the model in certain ways that will be explored in more detail in Chapter 4.

The agent behavior is considered to happen at three levels: an individual level, a group level and a crowd level. At the individual level, an agent generally acts using his experience. If experience can't help, they act rationally within the limits of their knowledge. If they are too stressed, then they start acting on *instinct*. Instinct refers to competitive behavior where pushing, jumping out of windows and

fleeing towards very crowded, blocked exits take place. At the interaction level, a social identity based model is used. Each agent has a social identity and each social identity has a set of actions associated with it. Depending on the situation, the actions taken by an agent are determined by the social identity. Pan assumes that during an extremely stressful situation people forget about their social roles. Personal spaces are defined for each individual which they always try to maintain. If this space is violated, they get agitated and stressed and will eventually lead to *non-adaptive behavior*. The model also assumes a strong herding behavior, where the first reactors influence the reaction of the rest. At the crowd level, the three factors that influence the behavior of an agent are crowd density, environment constraints and perceived emotion and tension. The first two factors are standard in most crowd simulation environments. The third factor is interesting; it states how the individual's perception of a system as stressful determines how stressed he is, not just what happens to him. So even in a non emergency situation, *non-adaptive behavior* might occur if a false alarm or some other stressful situation arises. Queuing, competitive, leader following, altruistic and herding behavior are modeled as resulting from the interaction of these different factors. A decision tree is used to simulate this behavior and to show how each of these behaviors is caused as a result of the aforementioned three levels. The actual implementation of the behavior at an individual level consists of 5 factors: familiarity (memory), decision making type (the intrinsic factors determined decision tree), urge to exit, stress threshold type and herding factor.

The movement engine has a simple collision detection algorithm and controls basic movements. Basic movement on stairways are also modeled.

The greatest strength of this model is the way in which the higher level behavior of agents are determined and described using a few fundamental lower level behaviors. It also has a more detailed visual perception system and movement system. However, it does not model pre-evacuation behavior or affiliative behavior of any sort. Like some of the other models, there is an undue prominence of non-adaptive behavior during evacuation.

2.3.4 Summary

In this section, the different types of computational models of fires ranging from network based models to smart environments and agent based models were introduced. Seven existing models that model the key features highlighted in Section 2.2.4 were also presented and analyzed.

- Pires's model [31] was one of the earliest to consider pre-evacuation behavior. However, the model abstracted away details of the pre-evacuation behavior as probabilities making it more difficult to study and analyze its effects.
- Still's thesis [3] highlighted the importance of data. He also, importantly, high-

lighted the fact it is not necessary to have a complicated model to simulate complex systems. Rather, simple rules may be able to produce the required complex behavior. Models should, where possible, follow Occam's razor, i.e. the simplest rules should be used. This was exemplified by Pan's MASSEgress model [98] which produced higher level herding, following and competitive behavior by using simple lower level rules.

- Pelechano's work [38] in combining a dedicated behavior model (PMFServ) with their existing crowd simulation model (MACES) demonstrated the importance of inter-disciplinary collaboration. By integrating an existing detailed emotion model to improve an existing computational model, they demonstrated the strengths and limitations of this approach.
- The Collective Panic Behavior Model [41] demonstrated how a theory of human behavior can be computationally modeled without having to make any abstractions.
- The EXODUS model [95] is one of the most detailed and mature models today. Hence, it is an excellent source of information on egress modelling and simulation.
- ESCAPES model [29] is significant for the importance given to modelling affiliative behavior and the spread of information.
- Pan's MASSEgress model [40] despite not having the behavioral complexity of ESCAPES is notable for the structured approach to modelling agent behavior and its ability to produce complex high level behavior using simple low level rules.

2.4 Summary of Literature Review

This chapter first introduced the complexity of the problem of modelling crowd evacuation and its multidisciplinary nature. In Section 2.2, the current theories on crowd behavior during an evacuation were presented. The salient features of human behavior during fire evacuation were identified from this and summarized in Section 2.2.4. In the following section, the different approaches to modelling crowds was introduced and some of the more significant models were presented in detail. Their strengths and shortcomings were summarized in Section 2.3.4.

This chapter has provided a broad overview of existing computational models of egress and briefly examined some of their strengths and weaknesses. Some of the more popular computational models were analysed in detail. Despite the number and variety of models, several of the salient features of human behavior during egress identified in Section 2.2 are generally ignored in existing systems. In order to address

this mismatch in the current understanding of human behavior during egress and the state-of-the-art in crowd modelling and simulation, the shortcomings in existing models need to be systematically identified and addressed. When perceived as a single problem of *modeling behavior accurately*, this task seems intractable due to the range of phenomena like pre-evacuation behavior and the impact of limited information processing capacity that need to be considered. To counter this problem, in the next chapter, the building blocks that make up a model of evacuee behavior during egress is identified from the literature on crowd behavior and its modelling. Subsequently, approaches taken by existing computational models to each of these parts of egress behavior are analyzed in detail and compared against the current understanding of these processes to identify research questions and solve them.

Chapter 3

The Building Blocks of a Behavior Model for Egress Simulation

In the previous chapter, the salient features of the behavior of a crowd engaging in egress were introduced. It is understood that people don't immediately exit a building on hearing a fire alarm or seeing smoke. The occupant gets only an inkling about the danger that is possible. In both these cases there is not enough information for him to get scared and make the decision to egress. If the cues are interesting enough, he then embarks on investigating to gather more information. On realizing that the situation requires some action, he forms a plan to evacuate either alone or, if possible, with his primary group members. In forming this plan of evacuation, unless he is trained for egress or very familiar with the environment, it is unlikely that he will have knowledge of all exits. As a result, he is most likely to just move to the nearest *known* exit. An overview of this process of human evacuation is illustrated in Fig. 3.1. In an agent based simulation, inter-agent variations in the parameters like the investigation and time thresholds and subjective knowledge would result in agents producing different evacuation patterns despite having the same underlying process.

As highlighted in the previous chapter, several parts of this egress process like the pre-evacuation phase and the effect of subjective spatial knowledge are not considered in existing computational models. In order to make the problem of modelling crowd egress behavior accurately more tractable, a small set of core subprocesses can be identified. These subprocesses, introduced in Section 3.1, and their interaction are necessary and sufficient to produce the egress process outlined above. Section 3.2 introduces some of the shortcomings of existing approaches to modelling each of these sub processes. Section 3.3 explains how these shortcomings are addressed and overviews the key contributions of this thesis.

Figure 3.1: The process of egress: This state diagram shows the different phases of behavior of a person engaging in egress and the triggers that cause phase changes

3.1 The Building Blocks of the Egress Process

The individual behavior of an evacuee during emergency egress was summarized in Figure 3.1 from the literature discussed in Chapter 2. This section outlines the building blocks of a system that simulates crowd egress during an emergency event. The first step in simulating the behavior of any human being engaging in an activity is modelling his *perception* of the environment. Once perceived, the information gained from perception is used to *identify events* and trigger key actions. In the case of an egress simulation, this would be evacuation (or one of the processes leading up to it). Following this, the person tries to form an evacuation plan using the *knowledge* he has of the layout of the environment and also the information that he is currently perceiving. Once such a plan is formed, a *path is planned* towards this point and he moves along this path towards his goal. Thus the process of egress simulation can be said to consist of four major building blocks: perception, event identification, spatial knowledge utilization and navigation. The following sections define and explain each of these building blocks and their significance in the evacuation process.

3.1.1 Perception

Perception refers to the process by which an environment is observed by a person. For an agent, this implies that it should extract features from the environment. These features are called percepts [99]. Each percept gives the evacuee certain information

about the environment. At the simplest level, percepts provide the agent information about the location and the environment around it and enables it to move naturally around the environment. In some cases, something out of the ordinary like a ringing fire alarm or smoke or even people running away from a location, might be perceived. Communication between agents can also be considered to be a mode of perception since it is a process by which additional information is gained by the agents. The perception process is the main interface of the agent with the environment through which it gets all the information needed to act. In order to act on information perceived, it is essential for the agent to be able to identify events that require some action.

3.1.2 Event identification

In some cases, evacuees might perceive something out of the ordinary like a ringing fire alarm or smoke or even people running away from a location. If these special percepts, known as *cues*, are intriguing enough, then the occupant will set about investigating the environment in search of more cues that would help make a decision on a plan of action. Cues are certain changes in the environment that indicate that something is wrong or different from normal [22]. Thus, cues help people identify an event and a need for an action. During emergency egress, on perceiving *sufficient* cues from the environment, an evacuee stops whatever task he's doing and initiates a process of investigation. Once enough evidence is gathered to convince the evacuee of danger, evacuation starts. If enough information is not obtained then the evacuee will likely resume their previous task. Thus identification of events is a key component of the evacuation process that is especially important for modelling pre-evacuation behavior.

In summary, event identification is the process by which the evacuee interprets perceived information and decides to proceed to the next phase of evacuation. Executing the action for a particular phase generally requires the evacuee to proceed to a particular location in the environment. This is done using the subjective knowledge that he has of the environment.

3.1.3 Spatial knowledge utilization

On recognizing that there is an event to escape from, the evacuee moves towards the closest known exit. Each evacuee holds his own personal version of the map of the environment which is used for evacuation and movement in general. This personal map is called his cognitive map [100]. In most cases, this map is not complete. In a scenario where all known exits are blocked or if the evacuee does not know any exit, he is likely to engage in some sort of exploration behavior which helps increase his knowledge and improve his cognitive map.

The process of knowledge modelling encompasses the method used for modelling the cognitive map and how it is obtained by the evacuee and his actions

when insufficient knowledge is available to act. Modelling evacuee-specific *knowledge* is necessary to model the heterogeneity in behavior found in real life emergency evacuations. The key function of the spatial knowledge model for the agent is to provide it with a target to move to complete the action required based on the event identified. The actual movement towards this point is controlled by the navigation system.

3.1.4 Navigation

Navigation is defined as the process or activity of accurately ascertaining one's position and planning and following a route. In the context of egress simulation, we define navigation as the process by which an evacuee first plans his route towards a *goal location* based on his cognitive map and then moves along this route towards the goal. Thus, navigation consists of 2 distinct processes: planning a route and following the route. The former is referred to as path planning and the latter is called motion planning.

During emergency egress, once evacuation has started, the agent decides that it must proceed to a particular exit, say exit D. Given the current location and this goal, the path planning system recommends a sequence of locations, or waypoints, through which the agent should proceed to get to this goal. Understanding the process of path planning in agent based simulations is easier if it is divided into two parts: A higher level path finder that finds abstract logical waypoints towards the goal, i.e. go through room A to room B to room C and take Exit D; and a lower level mechanism that translates these logical waypoints to physical locations on the map that the agent can pass through to reach its goal. Fig. 3.2, which gives a mathematical overview of the navigation process, refers to the former as “Logical Waypoint Determination” and the latter as “Current Waypoint Path Determination”. Once these physical waypoints are determined the actual movement towards these locations are determined by the motion planning process.

Motion planning is a term borrowed from robotics which originally means detailing a task into discrete motions. In the context of crowd simulation, we use the term motion planning to refer to the task of finding a collision free velocity to get from the current point to the next waypoint in the planned path. The motion planning layer ensures that the agent manages to reach its next way point without colliding with other agents.

Assuming that the cognitive map is modelled as a graph with the edges E representing *areas* or rooms in the environment, Figure 4.1 gives a mathematical model of how navigation is generally modelled.

The goal position ($g^{x,y}$) is first received by the *Logical Waypoint Determination* (WD) system which then uses a route finding algorithm like A-star or Dijkstra's to find a set of logical waypoints (E_i) that will lead the agent from its current position($p^{x,y}$) to the goal position ($g^{x,y}$). The logical waypoints in the path

Figure 3.2: Mathematical description of the working of a Navigation System

so planned will be a list of links from the agent's cognitive map ($Env_{cognitive}$). This implies two things: Firstly, it looks at the environment at the room/ area level. Secondly, the agent plans a path according to his personal cognitive map hence there is no guarantee that the path planned by two agents from the same point will be exactly the same. This step is called a waypoint determination step because it returns a set of logical waypoints to be used by the agent to reach its goal. This set of logical waypoints is then passed to the next level of navigation, i.e. *the Current Waypoint Path Determination* step.

The *Current Waypoint Path Determination* (PD) differs from the higher layer in that it takes into consideration dynamic obstacles i.e. other agents, as well as the static obstacles to determine a collision free path from the current location to the farthest visible waypoint. The logical waypoints are first converted to a set of concrete waypoints ($WP^{x,y}$) which are actual locations on the map. Using the set of perceived obstacles (O_{perc}) a set of intermediate waypoints ($IWP^{x,y}$) that would enable a collision free path to the farthest visible concrete waypoint ($wp_j^{x,y}$) is calculated and passed to the next level. This level can also be referred to as the *strategic planning step* since it tries to model how a person strategically moves from one point to another while ensuring that it minimizes collisions with others. Hu's pattern based

Table 3.1: The building blocks of human behavior during egress

Building Block	Definition	Purpose
Perception	The process of gathering information about the environment.	Learn about the environment and observe events.
Event Identification	The process by which the evacuee analyses and aggregates cues and identifies an event.	Change from one phase to another based on perception and internal state.
Knowledge Modellings	The process of using stored knowledge to formulate a plan for evacuation.	Determining a goal based on the current spatial cognitive map.
Navigation	The routing and movement process.	Handles movement towards the evacuee's current goal.

approach [101] and ClearPath by Guy et al. [102] are examples of models for this layer.

The preferred velocity of the agent is then set to the velocity that would lead it to the next intermediate concrete waypoint. Following this, the Motion Planning System like RVO [103, 104] or Social force [10] based methods determine the actual movement of the agent for the next time step of the simulation based on this preferred velocity. At each time step (t_i), the motion planning system takes a preferred velocity ($\vec{v}_{pref}$) and set of obstacles (O_{perc}) (both static and dynamic) as input and outputs a possible velocity ($\vec{v}_{possible}$) that the agent can use to ensure that collisions do not occur for the next N seconds. The two important points of difference from the previous layer is that it has some noticeable effect only when a collision is imminent and that it is very short term (limited to a few time steps) as compared to PD.

To summarize, the navigation system receives a goal that is passed to the highest level which determines a high level route through the different areas in the environment to the goal; this waypoint is passed to the next level which determines a path towards the current waypoint that avoids dynamic obstacles and calculates the velocity of the agent in order to get to the waypoint while avoiding collisions for the next time step of the simulation.

In summary, a perception system along with event identification, a knowledge model and a navigation system can together be used to produce the entire process of a person evacuating from a building. The function of each of these building blocks is summarized in Table 3.1. In the following sections, I first identify shortcomings in existing models of these processes and then explain the contributions of this thesis to those areas.

3.2 Shortcomings in Existing Models

The previous section identified the core building blocks of an agent based model of emergency egress. In Chapter 2, the salient features of human behavior during egress were identified and summarized in Section 2.2.4. In this section, this review is used to identify key research questions in each of the core building blocks, where either new modelling approaches or new understanding is necessary.

3.2.1 Perception: the lack of a realistic approach

Perception is the process by which agents perceive and gather information about the world around them. It is thus a crucial part of an agent based simulation of crowds. However, in most models, it is standard practice to consider a simple circular or elliptical sensor range. All agents and objects within that sensor range are perceived by the agent. One of the rare exceptions to this is the MASSegress model [40] that makes use of a ray tracing algorithm to model perception. The assumptions that an agent can perceive as many objects as are there in the sensor range without any internal limit on capacity [25] and that perception is simply a visual process are both simplistic. Another related consequence of the limited information processing capacity of humans is that the human brain tries to chunk similar information to improve its efficiency. Modelling this limited information processing capacity and chunking can produce significant improvements in the realism of crowd simulations. This is demonstrated in Chapter 4 which introduces an Information Processing Based Perception model which takes this into account.

3.2.2 Event identification: modelling pre-evacuation behavior

Several studies [35] have shown that humans do not evacuate immediately on hearing a fire alarm. This delay between hearing the alarm and actual evacuation has even resulted in deaths [28]. This can have a significant impact on the time to evacuation and the general efficiency of evacuation. However, despite several studies highlighting this, pre-evacuation behavior is hardly ever modeled in existing crowd simulations. In the few instances where it is considered [31, 105], the models are rather simplistic.

Chapter 5 introduces a model for event identification that enables the modelling of pre-evacuation behavior and studying of its effects. Through simulation of certain scenarios, the importance and usefulness of modelling pre-evacuation behavior is demonstrated.

3.2.3 Knowledge modelling: the effect of partial knowledge

On deciding to evacuate an evacuee plans a path towards an exit. Ideally, this plan would be to move towards the closest exit that is known. In most cases, however, say

a shopping mall or a hospital, it is highly likely that the majority of the occupants do not know emergency exit locations or the closest exits. How evacuees explore and form their memory in such a situation of partial knowledge is still an open problem.

Existing simulations either assume complete knowledge or have a simplistic memory model where a single visit is remembered forever. It is unlikely that evacuees have such eidetic memories. Also, it is common to simply take a depth first approach to exploration in case of incomplete knowledge. However, there is no proof whether this is the actual way in which evacuees explore.

Chapter 6 first introduces a game based methodology for studying exploration of indoor environments and how human memory works during exploration. By comparing against a pure random walker, it is shown that humans are in fact more effective explorers than a random walker but not as efficient as is portrayed in existing computational models of egress. The study also revealed some interesting patterns in the strategies used for exploring complex indoor environments.

3.2.4 Navigation: analyzing motion planning systems

The motion planning system used in an agent based simulation of egress is a key determinant of the kind of behavior produced by the model. The motion planning system used can have a great impact on the patterns of motion produced and the dynamics of the crowd. There are several different models of motion planning that have been proposed that sometimes work in very different ways. Each model has its strengths and experiments are generally directed towards demonstrating these strengths. However, despite the abundance of models, there aren't any standard methods for comparing different models against each other or against real world data. In Chapter 7 of this thesis, existing motion planning systems are quantitatively compared and evaluated and a metric that can be used for differentiating existing motion planning systems is determined. These motion planning systems are subsequently compared against real world data.

3.3 Contributions of the Remaining Chapters

Following is a summary of the contributions of the remaining chapters:

- **Chapter 4:** This chapter presents an Information Processing Based model of perception which was presented at the Cyberworlds 2011 Conference [106] and further explored in the extended version published in the Transactions on Computational Science [11]. This chapter additionally contains some validation of the model against real world experimental data.
- **Chapter 5:** This chapter presents a method for modelling pre-evacuation behavior in emergency egress simulation. The model was presented at the Pedes-

trian and Evacuation Dynamics Conference in 2012 and published in 2014 [107].

- **Chapter 6:** A game based methodology for studying the effect of memory on exploration of complex indoor environments and different exploration strategies used is presented in this chapter. The work done in this chapter has been submitted for review to Behavior Research Methods.
- **Chapter 7:** The final contribution of this thesis is a methodology for quantitatively comparing different motion planning systems which was also used to compare some popular models against real world data. This was published in European Physical Journal B [108].

3.4 Summary

This chapter first summarized the behavior of an evacuee during evacuation from the literature in the previous chapter. A comparison of this egress process against existing computational models reveals several shortcomings in existing approaches to behavior modeling in egress simulations. In order to identify and address these shortcomings systematically, the four essential building blocks for an agent based simulation of the emergency egress process was introduced. Subsequently, an overview of the shortcomings of existing approaches to modelling each of these building blocks were identified and the contributions of the remaining chapters was introduced. The following chapters of the thesis present these contributions starting with a more realistic model of perception.

Chapter 4

An Information Processing Based Approach to Modelling Perception

A core component of an agent based crowd simulation is the model of movement that is used. Naturally, there has been a lot of work done on developing realistic movement models. As with any system, the effectiveness of a movement model depends on its input; in this case it is the model of perception that is used. However, this is also an aspect of motion planning that is generally not given much consideration. In this chapter, we demonstrate how a more realistic model of perception can improve results produced by existing motion planning systems. This is done through a model of perception that takes into consideration the limitations of human perception.

As discussed in the previous chapters, Agent Based Simulation has recently become the preferred approach to simulating crowds because of the level of detail in which crowds can be modeled. Agent-Based Models (ABMs) consist of large-numbers of heterogeneous, autonomous entities inhabiting a spatially explicit, partially observable environment; macro level dynamics are said to emerge through the asynchronous interactions among these entities [87, 88]. Each of these individual entities will iterate through a sense-think-act cycle [90], where agents obtain information from their environment through *sensing*, make a decision through *thinking* and finally carry out their decision by *acting*. In many application areas in which ABMs have been applied, including crowd simulation, the emphasis is generally on describing thought processes accurately via rules. However, sensing is a critical aspect in the modelling process and can greatly impact both the individual and emergent properties of the system.

The terms perception and sensing are often used interchangeably in the simulation literature. For clarity in explanation, the term *perception* is used here to refer to the complete process of obtaining a set of (possibly filtered) *percepts* [99] from the environment. *Sensing*, on the other hand, is defined as the process of obtaining raw information from the environment; in this definition, and in this model, sensing is a part of perception.

The sense-think-act cycle is the process by which humans get information from the environment, process this information to make a decision and, finally, act based on the decision made. Rather than what a human sees, hears or smells, what is more important is what he can mentally process. In fact, the entire human perception system can simply be considered to be an information processing entity. This chapter presents a perception system based on the idea that perception is the process of gathering information from the environment. This is called the *Information Based Perception*(IBP) system.

Miller's work [25] on human cognition revealed two important characteristics of how the human brain processes information: 1. Humans constantly group together similar data into *chunks* of information. 2. At any given time, a human can only process a limited amount of information.

For IBP, the assumption is made that this limited capacity results in humans being attracted towards certain kinds of information, e.g. a bright light or a celebrity; this, in turn, results in other information in the environment being unnoticed. By organizing information into chunks, humans are able to use their limited information processing capability more efficiently. This ability can manifest itself in different ways. It can be reasonably assumed that during motion planning, humans will process a group of people coming towards them as a single obstacle rather than many individuals. This grouping not only helps the person make use of his limited information processing capacity more efficiently, it also helps him conform to social norms that instruct him that walking through a group of interacting people would be rude.

This chapter explains and illustrates the working and usefulness of an Information Based Perception system for agents. Its viability is demonstrated through the implementation of a simple moving agent and by incorporating information based processing into the agent's motion planning system. Besides being one of the major components of Agent Based Crowd Simulation models, motion planning is also a process in which the effects of using a new perception system can be observed easily. The experiments towards the end of this chapter illustrate the significant effects that a modified sensing and perception system can have on an existing *motion planning* algorithm. The results produced are also compared against real world experiments in an attempt to validate the model.

The remainder of this chapter is organized as follows: Section 4.1 gives some background on how humans perceive the world around them; following this, Section 4.2 presents a brief overview of some of the existing work in motion planning; the IBP model itself is introduced in Section 4.3; Section 4.4 presents the work done in visual, experimental and quantitative validation; finally, Section 4.5 concludes this chapter and gives an overview of the work that needs to be done.¹

¹This chapter was presented at the Cyberworlds 2011 Conference [106] and an extended version was published in the Transactions on Computational Science [11]. This chapter additionally contains

4.1 Limits of Human Perception

In 1953, Hochberg and McAllister [109] proposed the idea that humans try to group together similar information so that information can be encoded in the simplest possible format. They called this *the simplicity principle*. This theory was further extended by Miller [25] who proposed that, at any given time, humans can only process a limited amount of information. He explained this as the human short term working memory having a limited capacity. To enable humans to store more information in this limited storage space, humans “chunk” together similar information. It was originally proposed that the short term working memory could hold 7 ± 2 chunks. Cowan [110] has argued that this limit is actually 4 ± 1 for most humans. Thus, even though a person’s 5 senses are giving him a constant stream of information about the world, limitations of human short term working memory force the person to act on the basis of only a fraction of this received information.

With regard to visual perception, some studies [111–113] have shown that humans only pay attention to certain salient features in the objects that they see. This results in them not noticing changes in items that are not of interest to them. O’Reagan et al. [112] classified elements as either central interest or marginal interest elements and showed that the internal representation of the visual world is rather sparse and essentially contains only central interest information and not information of objects of marginal interest. The world that we perceive around us is a combination of this sparse visual world along with the information in our working short term memory received from our other senses.

Based on these studies, for our model, we make the reasonable assumption that the human brain uses some mechanism to determine the significance of a particular *raw percept*. And the short term working memory stores the most significant information in its limited capacity. It is important to realize that this significance determination is done for *all* information received, regardless of the source. We call this significance, the *amount of information*.

The idea of considering the human perception system as an information processing system is not unprecedented. Broadbent [114] extensively discussed the idea of using information theory for modelling human perception. Various studies were presented that indicate that humans have an upper bound on their capacity for holding information for perception. For a single dimension, this limit is roughly estimated to be about 5-6 percepts. For more than one dimension, the number of discernible alternatives is larger but not as large as would be expected if each dimension was completely independent.

The idea of humans being able to process only a limited amount of information is not new to computer animation either. Hill [115] was one of the first to introduce the importance of cognition in sensing. Courty [116] used a saliency map

some validation of the model against real world experimental data.

based approach and Kim et al. [117] used cost-benefit analysis in a decision theory based approach to determining the interest points. Grillon and Thallman [118] automated this process of interest point determination. They used criteria like proximity, relative speed, relative orientation and periphery to determine the interestingness of various features.

The majority of existing perception systems, consider perception to be only visual perception. Even in more detailed crowd simulation systems like LE-GION [3] and MASSEgress [40] perception is implemented to aid movement by detecting other obstacles and goals to enable planning a path towards the goal and to provide a collision free motion. The IBP System implemented and demonstrated in this chapter is similarly limited, i.e. the IBP is modeled in the context of collision avoidance. The only information which the agents perceive are static obstacles, like walls, and dynamic obstacles, i.e. other agents or groups of agents. However, in theory, the concept of an information processing based perception system can be extended to include factors like perception of fire, smoke or other more complex *cues*. Cues are certain changes in the environment that indicate that something is wrong or different from normal [22] and modelling of cue perception is an important part of egress simulation. This will be discussed in more detail in Chapter 5.

In the present model, the idea is not to model all the complexities of human perception and visual cognition; rather a perception model for producing more realistic movement in agent based simulation of crowds is presented. The next section gives a brief review of motion planning in existing crowd simulations.

4.2 A Brief Introduction to Motion Planning Systems

Navigation is defined as the process or activity of accurately ascertaining one's position and planning and following a route. Thus we use the term *navigation* to refer to the complete process of how a person moves from one point to another. Navigation itself can be broadly divided into 3 (or 4 parts) as shown in Fig. 4.1. In this section, a brief overview of *motion planning* is given. A more detailed description and analysis is given in Chapter 7. In a simulation, motion planning ensures that the simulated human avoids collisions during movement.

There are several different approaches to motion planning. For example, Okazaki and Matsushita [119] used a magnetism based approach to motion with all agents having the same pole so that they repel each other and with the goal having the opposite pole. Klein and Köster [80] used a similar approach of using coulombic charges instead of magnetic poles; they assigned positive charges to goals and negative charges to obstacles and agents. The most popular force based model, however, is Helbing's Social Force Model which is discussed in more detail later. Individual based

Figure 4.1: Navigation architecture

motion planning was first introduced in the flocking model of Craig Reynolds [7]. More recently, there have been several velocity-based approaches to motion planning, such as the synthetic vision based model [120] and the Reciprocal Velocity Obstacle (RVO) Model [103, 121, 122]. Of the many models, variants of RVO and Social Force [10] are the most popular.

Helbing's social forces model was first introduced in 1995 [10]. In this model, each agent is modeled as a particle that has multiple forces acting on it. Repulsive forces help in collision avoidance and attractive forces model goal directed and grouping behavior. Over the years, this model has been extended and combined with other higher level behavior models. For example, in [123] more complicated group movement was modeled with an underlying social forces model for collision avoidance. In his thesis, Still [3] criticized the heavily mathematical approach which, according to him, is too complicated to be the natural way in which humans try to avoid crowds.

Another ABM that is increasingly becoming popular for collision avoidance is based on the idea of using the relative motion of objects to determine their time to collision. A velocity is then selected which maximizes this time. This algorithm, based on RVO was first extended for use with multi agent systems in [124]. Since then there have been several modifications and improvements to the system but the underlying algorithm still remained the same. CLEARPATH [102] which mathematically optimized RVO was the first to introduce a change in the underlying algorithm. Guy et al. [103] introduced an entirely new approach to RVO that was

based on computational geometry and linear programming. This method further improved the efficiency and smoothness of the system and was called *RVO2*. In another article, Guy et al. [12] introduced a personal space factor and an observation delay making the algorithm more appropriate for virtual humans.

Guy et al. [12] introduced an extension to RVO in the form of a higher level navigation based on the principle of least effort. While it is obvious that rational humans would prefer taking the path of least effort, as was explained in Section 4.1, humans do not have perfect knowledge or perfect calculation. Also, it is arguable whether humans are always rational enough to choose least effort as their goal.

There are a number of existing motion planning methods that can effectively and efficiently calculate trajectories that avoid all collisions for agents, even in relatively dense environments. For robots and computer games, this might be the ideal goal: perfect, smooth and efficient motion. However, for applications like simulation of emergency evacuation the goal is obtaining realistic motion and not smooth and efficient motion. While humans thrive to be mechanically efficient, this is hardly always the case. There exist, among other things, social norms and limits to mental processing capabilities that prevent individuals from following their ideal preferred path. We believe that a more realistic perception model like IBP, which takes into consideration human inadequacies and limitations, can help produce more naturalistic results [80] with existing motion planning systems.

In this chapter two additions to general motion planning algorithms are proposed: 1. Group sensing for motion planning which results in agents avoiding clusters of other agents when choosing their collision free path. 2. Filtering of percepts based on the amount of information provided to model limited information processing capabilities of human beings.

Guy et al. [12] clustered very distant objects into KD-trees to reduce computational cost. While this might sound similar to the idea that is suggested in this chapter, there are two fundamental reasons why this is different from the algorithm presented here: Firstly, the present model uses multiple levels of clustering which will be explained in more detail in Section 4.3.1. Secondly, the motivation and hence design is significantly different: clustering in IBP is used as a reflection of how agents perceive their environment and not an optimization for collision avoidance.

In the following section, a method that will emulate how humans perceive groups whenever possible and a system in which the agents avoid these groups rather than individuals is proposed. This has been done using the Evolving Clustering Method (ECM) [125] and computational geometry based RVO2 [126]. But our approach can, in principle, use almost any clustering and collision avoidance algorithms.

Figure 4.2: An agent perceives and then acts

4.3 The Information Based Perception Model

This section explains the Information Based Perception System. Figure 4.2 illustrates how motion planning works in an agent in terms of a sense-think-act cycle. An agent's perception can be described by a function $f : Env \rightarrow p^*$, where p^* is the set of percepts. Each percept p is then processed by the agent in its decision making process, which in turn will determine an appropriate action for collision avoidance. In our case, the motion planning module is passed a set of percepts which consists of neighboring agents and static obstacles which it processes to find the most appropriate velocity for reaching the goal. Typically, this list of neighbors is a set of agents within some cone of vision or some distance away from the agent. In the proposed IBP, a modification to this traditional perception procedure is proposed such that it takes place in three phases: clustering, sensing and filtering. As discussed in the introduction, this is done in order to simulate the human tendency to group together similar information in order to make more efficient use of their limited short term memory capacity. The effectiveness of the proposed method in producing more realistic results is demonstrated in Section 4.4. Figure 4.3 gives an overview of the process that is detailed in the following sections.

4.3.1 Clustering

Central to our information based perception system is the definition of *information units*. In traditional crowd simulation each individual agent or obstacle is considered as a percept, i.e. as an entity which should be processed by the motion planning system. In our approach, we assume that percepts can be either individuals or groups of other pedestrians. Whether an individual considers a group or individual is related to the *coherence* of the group and also the distance of the perceiving agent from the group. In order to achieve this, we perform a global clustering across the entire environment of agents. However, if an agent perceives a very big group close to it, then the agent may not have the time to go around this obstacle. To counter this, we

Figure 4.3: Perception in agents takes place through three stages: (1) Clustering is done at a global level. The dotted line indicates this separation. (2) Sensing is the process by which the Agents perceive only a subset of this. (3) Filtering further reduces the size of this list and models human visual cognition

create n_l layers within the environment, each layer identifies and stores groups of a particular maximum size, with higher layer numbers storing groups of bigger size. As explained in more detail in Section 4.3.2, whether an agent senses a bigger group or a number of smaller groups is determined based on the distance. The criteria which determines what actually constitutes a group is itself unknown and probably highly dependent on the individual. We make the assumption that only the proximity of the individuals to one another determines whether a collection of people is perceived as a *group*.

For reasons of efficiency we perform a single clustering (for each level) for all agents at every time-step, the consequence is that we are implicitly assuming all agents have the same notion of what constitutes a group. In reality this assumption may be too strong, different people may have different criteria for what they perceive as groups.

While there are various clustering techniques that could be used for grouping agents, we chose to use ECM [125] because: 1. It does not require the number of clusters to be predefined and 2. It can restrict the maximum radius of a cluster. It is also important to remember that this clustering is done dynamically at *each step* and not as a one time calculation of groups.

First, the number of clustering layers is decided. More layers would

result in the agent being able to perceive larger groups. However, groups that are too large are likely to be outside the perception range of the agent and thus unnecessary. Through experimentation we found two layers to be suitable. In Fig. 4.4, we illustrate information based perception using two layers. The algorithm starts by initializing a single agent as the first cluster, the maximum clustering radius for layer i , r_{max}^i is fixed (Equations 4.3 and 4.4). Each subsequent agent is then compared with every existing cluster to assess its suitability for addition to that cluster. Suitability is determined by the distance of the agent from the cluster. If the agent lies within an existing cluster, it is simply added to that cluster without updating either the radius or the cluster center. Otherwise, the cluster whose center is closest to the agent is determined. If the agent can be added to this cluster, without exceeding the maximum allowed radius for the cluster, then the agent is added to the cluster and the cluster's radius, center and velocity are updated. On the other hand, if adding the agent violates the maximum radius criteria, then a new cluster is created at the location of the agent.

Once this process is completed for layer i , the process is repeated for layer $i+1$ until the clusters for all the layers are determined. This process is illustrated figuratively in Fig. 4.3. The clustering function for layer i , cf_i allocates one and only one cluster for each agent in each layer. This can be represented mathematically as shown below:

$$\forall a_k \in A \exists j \in [1, m] \quad cf_i : a_k \rightarrow C_{ij} \text{ where } 1 \leq m \leq n \quad (4.1)$$

$$\forall a_j \in A \quad C_{0j} = a_j \quad (4.2)$$

$$r_{max}^1 = 2\alpha * r_a \quad (4.3)$$

$$\forall i \geq 2 \quad r_{max}^i = 2\alpha * r_{max}^{i-1} \quad (4.4)$$

Here r_a is the average radius of an agent² in A , which is the set of all agents; C_{ij} indicates cluster j in layer i ; m is the number of clusters and n is the number of agents. α is a parameter that determines the size of clusters and the range of each region (Fig. 4.4). Through experimentation we found an α value of 2 to be the most suitable.

The ECM based clustering for each layer considers each agent exactly once, so the process has an asymptotic complexity of $O(n^2)$. At the end of each clustering, each agent belongs to a cluster. However, in the absence of nearby agents this might be a singleton cluster.

To correct certain undesirable behavior produced by ECM clustering, a modification was made to the algorithm. With large values of r_{max} , there is a chance that distant agents might be grouped into sparse clusters. To counter this problem, we define an *inspection area* as a circle of radius $2\alpha r_a$. If there are no agents within this inspection area, then the cluster is considered sparse and the cluster is removed.

²In the experiments at the end of this chapter, it is assumed that all agents have the same radius. Hence, the radius of every agent is the same as the average radius.

Figure 4.4: The figure illustrates how the opaque agent senses objects using 2 clustering layers. The bottom layer is the original environment and the two planes above show the two clustering layers. Clusters in layer 2 are generally bigger than in layer 1. Solid lined circles indicate the normal agents and the clustered agents. The dashed line demarcates the regions of perception.

The sparseness check is done five times: First with the inspection area centered at the center of the cluster; and subsequently with the inspection areas centered at a distance equal to half the distance from the center of the cluster along each of the coordinate axes.

4.3.2 Sensing

Once the agents have been clustered, the next step is to make use of these clusters for motion planning. As previously explained, existing motion planning algorithms need a list of nearby agents and obstacles to determine the most appropriate velocity. The sensing module of our proposed perception mechanism uses the set of n_l layers created in the clustering module. The list of things to avoid will now consist of agents, obstacles and groups of agents. This list of nearby objects is now calculated from the multiple clustering layers as shown in Fig. 4.4.

From each cluster layer (explained in Section 4.3.1) a *perception region* pr_i is defined for each agent. This region can be considered as a modification of the sensor range which is used in most ABM and is assumed to be a ring-shaped region. In the first region (pr_0), immediately surrounding the agent performing the sensing, the agent perceives other individual agents from the clustering layer 0. This region extends to a distance $r_{pr_0} = 2\alpha * r_a$ from the agent's current location. For each subsequent region, the ring shaped region of sensing is from the boundary of the

previous layer's region to the boundary of a circle of radius 2α times the radius of the preceding region. So for region pr_1 the agent perceives groups of maximum size r_{max}^1 , as long as the nearest edge of their minimum enclosing circle is within a distance d , such that $r_{pr_0} < d \leq r_{pr_1}$. The result is a list of obstacles which consists of clusters of various sizes and individual agents.

4.3.3 Filtering

As explained in Section 4.1, a human being does not cognitively process every single object or obstacle that is within its sensing range. In other words, an agent can only process a limited amount of information and the information that is processed will be that which is deemed most interesting or important to the agent. So each object in the list obtained from perception is assigned an interestingness score of between 0 and 1 (or greater than 1 for exceptional cases discussed later). During the sensing process each agent is given an information limit il_a , indicating the total amount of information that can be processed by the agent. This limit is a parameter that can then change as the stress level or other characteristics of the agent changes [26].

For now, it is assumed that interestingness of an object depends on two criteria: 1. The distance of the object from the agent. 2. The angle that the object currently forms with the direction of motion of the agent. A third factor indicating the innate interestingness of the object being perceived can also be used. This can be used as an aggregate to represent a lot of other properties related to interestingness. For example, an object's speed, color, action or something more subjective, i.e. it is of interest only to this agent because of certain properties of the agent. For e.g., for a thirsty agent, a water cooler would be interesting, whereas it is unlikely to catch the attention of someone else. A more exact definition of interestingness is not the focus of this chapter, but the general model should be able to adapt to more sophisticated definitions. However, a notion of interestingness is required to extend IBP to detect events and cues in the situation and environment.

Based on the two criteria, a score is given to each agent. A distance score of $S_{d_{immed}} > 1$ is given if the distance between two agents is less than or equal to zero. This is to ensure that in high density scenarios where a collision does occur, a collision recovery mechanism is forced on the objects regardless of what angle or how interesting the object is. Too large a value would mean that too few other obstacles would be considered leading to new collisions. Through experimentation, we found a value of $S_{d_{immed}} = 1.5$ produced suitable results. Beyond this, within a certain range, the distance score is fixed at 1.0. This is considered to approximate the normal sensing range of the agent and is fixed at $7m$. Beyond this, the distance score is assumed to fall exponentially. The following equation defines this function:

$$S_d = \max(\min(1.0, e^{\gamma/d} - k), 0.1) \quad (4.5)$$

Figure 4.5: This graph shows the variation of distance score with distance (in metres) used in experiments. A score of $S_{dimmed} > 1$ is given if a collision has already occurred. Through experimentation this was fixed at 1.5. The model assumes that within a certain distance agents are almost definitely perceived regardless of their distance. A distance score of 1 is assigned to objects in this range; the distance score beyond this decreases exponentially with distance.

γ and k are parameters which were fixed at 5.0 and 1.11 respectively to get a curve as in Fig. 4.5.

An angle score of 1.0 is given to all objects forming an angle of less than a_{min} with the agent's direction. For all agents that form an angle of more than a_{max} with the agent's direction, a score of $(1 - \beta)$ is given. For our experiments a β value of 0.9 was used and this is illustrated in Fig. 4.6. For all angles in between, the angle score linearly decreases to $(1 - \beta)$ from 1. This is assigned based on the following equation (Fig. 4.6). All angles are in radians:

$$S_{\theta} = 1.0 - (\beta * (a - a_{min}) / (a_{max} - a_{min})) \quad (4.6)$$

The final score for the object is calculated as the product of the S_{θ} and S_d (as long as the distance score is not 1.5). This list of objects is then sorted on the basis of the score that is determined. Objects are then taken from the head of this list in turn and added to the final list of perceived objects as long as the cumulative score of all the perceived objects does not exceed the information limit for the agent, il_a . All the remaining objects are dropped from the list of objects sensed and the final list of percepts p^* is obtained. In case two objects have the same score, the objects that are moving towards the perceiving agent are given precedence, subsequently closer objects are given preference.

This shortened neighbor list is passed as input to RVO2 [103] for cal-

Figure 4.6: This graph shows the variation of angle score with the angle(in radians) formed by the object with the agent used in experiments. For objects forming an angle of less than 70° (viewing angle 140° , a score of 1 is given. For objects forming an angle of up to 90° , the score linearly decreases to 0.1 which is the angle score for all remaining obstacles.

culating the velocity at each time step. Our hypothesis is that the 3-step perception process presented in this chapter provides an improvement in two ways: Firstly, there are fewer neighbors and hence, fewer constraints for a given sensor range. Secondly and more importantly, more natural results can be obtained as will be illustrated in Section 4.4.

4.4 Model Evaluation

This section evaluates the validity of the model in three ways. First, Section 4.4.1 presents a visual and quantitative demonstration of the usefulness of the proposed model. Following this, in Section 4.4.2, it is demonstrated that in standard scenarios the IBP enhanced RVO2 model produces results that are at least as good, if not better, than those obtained by using a traditional sensor range. Finally, Section 4.4.3 validates the model against real world experiments.

4.4.1 Model demonstration

The ideas introduced in Section 4.3 are used as the basis for visually validating different aspects of the proposed model. For quantitative validation of the model, two measures are used: *Effort Expended* and *decision inconvenience Cost*. In proposing their least effort based approach to motion planning [12], Guy et al. used a measure of effort expended to demonstrate the usefulness of their model. This effort was

Figure 4.7: Experiment 1: Group Based Perception

calculated as follows:

$$E = m \int (e_s + e_w |\vec{v}|^2) dt \quad ^3 \quad (4.7)$$

In this section, the same measure of effort is used to analyze and validate the proposed IBP model. For simplicity, all agents are taken to have the same average mass of 70 kg. However, this only measures the mechanical effort involved. To measure the amount of effort spent in decision making, an *decision inconvenience cost* measurement is introduced. The decision inconvenience cost is the number of time steps in which the agent chose a velocity other than its preferred velocity i.e., the number of times they had to avoid a collision.

Four different scenarios are considered to evaluate the overall performance. First, the effect that Group Based Perception can have on an agent moving through a crowd is demonstrated and the scenario is visually and quantitatively analyzed. Next, the effect of the multi-layered clustering on an agent moving towards a large group is similarly analyzed. Following this, the necessity of Group Based Perception (GBP) in modelling the information processing limits of human beings is shown. Finally, the importance and relevance of the information threshold is demonstrated by demonstrating the effect that it can have on an agent.

4.4.1.1 Group Based Perception

In this experiment the results of using RVO2 with a traditional simple circular sensor range against RVO2 with a Group Based Perception system is shown. The intention is

³ $e_s = 2.23 \frac{J}{Kgs}$ and $e_w = 1.26 \frac{Js}{Kgm^2}$ for an average human [127]

Table 4.1: Quantitative analysis of Group Based Perception

Agent Considered	Effort ($Jm^{-1} * 10^5$)		Decision Inconvenience Cost	
	Without GBP	With GBP	Without GBP	With GBP
Black Agent	71730	71726	120	148
All other agents (average)	1884	1880	14.28	6.56

to show the effect of perceiving agents as groups. The hypothesis is that by perceiving groups as obstacles the simulation will generate more natural motion. In Fig. 4.7, there is a single black agent moving towards the right, and a number of groups of red agents moving towards the left. The black trail shows the path that is taken by the black agent. It can be seen that in Fig. 4.7a where GBP was not used, the agent walks through other groups. Since RVO2 enforces each agent to do half the work to avoid collision, the agents within the group individually give way through its center for the oncoming agent to pass. However, as shown in Fig 4.7b, the GBP algorithm is capable of generating motion which avoids entire groups. The argument that this is more natural is based on the discussion in Section 4.1; due to social norms and the human tendency to group information together people generally try to move around an entire group rather than walking directly through a group.

An analysis of the effort expended and the decision inconvenience cost gives some interesting results. Since the simulation is executed for a given number of time steps, the effort expended is normalized with the progress towards the agent's goal. This is to avoid slow or stationary agents from being considered to be more efficient despite traveling a lesser distance. On comparing the normalized effort in the two scenarios of the black agent, it is found that despite having a much longer path, the GBP enabled agent expends slightly lesser (practically the same) amount of effort than the other. This is because the non-GBP agent has to slow down to wait for the other agents to give way before it can proceed and thus progresses less towards the goal.

The decision inconvenience cost comparison gives another interesting, though not surprising, result. The decision inconvenience cost to the black agent of using Group Based Perception is higher because of the more indirect path that it has to take. However, the average decision inconvenience caused to all the other agents is significantly lesser. This is consistent with the general human reluctance to decision inconvenience others. It also gives the interesting idea that even if the same amount of mechanical effort is expended in following two different paths, the amount of decision making required for each path might be significantly different.

Figure 4.8: Experiment 2: Effect of multi layered clustering

4.4.1.2 Effects of multi layered clustering

In this experiment, the simple scenario where a single (black) agent had to get past a big group of agents to get to its goal is studied. The same experiment was performed by keeping the agent at different distances from the group. The objective of this experiment is two-fold. Firstly, it demonstrates the importance and the working of the multi-layered clustering (Section 4.3.2) used. Secondly, it demonstrates that when agents are very close to each other, where RVO2 already performs well, the Group Based Perception does not interfere with RVO2's functioning.

To recap, the multiple layers are used to describe groups of varying size at varying ranges of perception. This means agents will perceive other agents as groups or individuals depending on the distance; as an agent moves towards a group it will start to perceive the group as individual agents.

When GBP isn't used, the path followed does not change significantly with distance. The agents in the path of the black agent, give way to the agent, and the black agent just proceeds straight through the center of the large group (Path I in Fig. 4.8a). In the last few cases (Paths II and III), the path is slightly different because the black agent does not have enough time to plan for a smooth, straight path

Figure 4.9: Experiment 2: Quantitative analysis

and hence there is a slight deviation. Also, similar to the experiment in Section 4.4.1.1 it is forced to slow down in the process.

The result produced when GBP is used is more varied. Four distinctly different paths (labeled IV, V, VI and VII in Fig. 4.8b) are produced based on how far the oncoming black agent is from the big group. At distances between 7-18m away from the center of the big cluster, the agent has enough time to perceive the group and avoid it completely (Path IV). At distances between 5-7m away, due to the size of the group, the agent gets too close to the group such that it then perceives the group as individuals. At this time (as described in Fig. 4.4) the agent performs motion planning on all the individual agents and as a consequence moves through the group, shown by path V. Path VI is obtained in a similar fashion; however, the black agent is too close to the group (4m away) to discern the effect of GBP. At distances closer than this (2-3m away), the path followed by the agent (Path VII) is exactly the same as that followed by the agent not using Group Based Perception (Path III). We argue that this type of flexibility in the perception of groups is critical to creating more natural behavior, humans will adapt what they perceive based on success or failure of their attempt to avoid larger groups.

Figures 4.9a and 4.9c show a comparison of the effort expended by the black agent and the average effort expended by all the remaining agents while using

a traditional sensor range and GBP. As in the previous experiment (Section 4.4.1.1) there is hardly any difference in the effort expended in both scenarios (except for a slight increase for path V). However, an interesting pattern can be observed in the decision inconvenience measurement (Figures 4.9b and 4.9d). Firstly, the decision inconvenience for the rest of the group, is always lesser when GBP is used and almost the same for the black agent when path IV is followed. However, when path V or VI is followed there is a spike in the decision inconvenience curve. This can be explained by considering the fact that in both path V and VI, the black agent changes its planned path suddenly and decides to go through the group, thus not only increasing its own decision inconvenience but also the decision inconvenience caused to others in the group who have to move to give way to the agent. Finally, when path VII is followed both the effort and decision inconvenience count are exactly the same as for path III.

4.4.1.3 Filtering necessitates Group Based Perception

In Section 4.3, the fact that humans have limited information processing capacity was explained. In this experiment, it is demonstrated that if a human being's limited information processing capability is to be modeled, it is necessary to use Group Based Perception. This is done by observing the simple scenario of an agent moving towards two groups of other agents (Fig. 4.10). When no information limit is imposed on the agent, and a normal circular sensor range is used, the agent, as expected, follows a nice straight path through the center of the group. However, when an information limit of $il_a = 4$ is imposed on the agent, the black agent, does not perceive all the individual agents in the group and as a result it is forced to reconsider its path mid-route. As a result, the irregular trail shown in Fig. 4.10b is obtained. However, in the same situation, when Group Based Perception is used, the agent smoothly avoids the whole group (Fig. 4.10e). In fact, this smooth path is obtained for as low a limit as $il_a = 1$ (Fig. 4.10d).

4.4.1.4 Effect of filtering of percept information

The final experiment (Fig. 4.11) demonstrates the effect of filtering, i.e. having limits on the information processing capabilities of the agents. The scenario consists of an agent moving towards a collection of individuals (moving towards the agent) followed by a group of agents behind the set of individuals. In the first case an information limit of $il_a = 5$ is set so that the agent is continually capable of perceiving a larger number of other agents and groups. In the second scenario a lower limit of $il_a = 3$ is used such that the agent isn't initially capable of perceiving the group behind the individuals. Figure 4.11c shows how agents perceive the cluster that is farther away, even when there is an immediate collision to avoid. Figure 4.11e shows that the agent manages to move around this group because it had a head start in planning - i.e. , it considered the group early when avoiding collisions. In the second scenario it

Figure 4.10: Experiment 3: The necessity of Group Based Perception

could process a maximum of 3 or 4 percepts at any given time because of the lower information limit. Due to this, as seen in Fig. 4.11b, the agent cannot see beyond the immediate obstacles in front and does not prepare in advance to avoid the larger group. Once the agent finally perceives this group, it is too late to move around this group as it perceives the group as individuals and then moves through the group as in Fig 4.11d.

This experiment illustrates how small differences in the information limit can generate different forms of behavior in the agents. Interestingly, the info limit of 3 and 5 correspond to Cowan's finding [110] that all humans can cognitively process only 3-5 chunks of information at any given time. Clearly the value of the limit is critical to behavior, it is also proposed that this limit will change with personal characteristics and the emotional state of the agents. In fact this varying limit of perception may be an important factor for collisions in actual crowds, this is especially relevant in emergency egress scenarios where stress and collisions are critically

Figure 4.11: Experiment 4: Effect of filtering of percept information

important to safety planning. We plan to attempt to quantify this information limit through experimentation in future work.

In this section, we demonstrated the capabilities and the usefulness of the Information Based Perception system in Motion Planning by analyzing the effect of group based perception, multi-layered clustering and the filtering of information. In the next section, we demonstrate that RVO2 when used with IBP produces results that are at least as good as with a traditional sensor range by analyzing its performance in standard scenarios.

4.4.2 Validation against standard scenarios

Several studies have identified standard scenarios against which motion planning systems have to be validated empirically. Hu [1] classified these scenarios based on their main motion axes (*mma*) into three main categories. Based on this categorization,

Figure 4.12: Corridor Experiment: Scenario with 1 main motion axis. With IBP, the lane formation is slightly more asymmetric

he identified that validation against scenarios with 1, 2 and greater than 2 *mmas* would be necessary to validate a model. In this section, we demonstrate that in these three standard scenarios, IBP enhanced RVO2 produces motion similar to traditional RVO2. For a more detailed discussion on the setup of these scenarios and the reason for this categorization, the reader is directed to Hu’s thesis [1].

4.4.2.1 Corridor (1 *mma*)

Hu et al. [1] identified a corridor with agents starting at either end and moving towards each other as the standard scenario in this category. Figure 4.12 illustrates the simulation results produced by RVO2 with a traditional simple circular sensor range against RVO2 with a Group Based Perception system. An empirical observation of the simulation results does not seem to indicate much difference between the two models.

Traditionally, in this scenario, it is expected that the pedestrian agents form lanes to handle the counter flow and prevent jamming [128]; this is seen in both models in Figure 4.12. Based on more detailed studies done in [128], Hu et al. suggested several qualitative measurements of the lane formation phenomenon. Firstly, they observed that lanes with a longer lifetime can more reliably reduce jams and disordered flow. As seen in the figure, both models simulated seem to produce lanes that last the length of the simulation. They also suggested that, due to the pedestrians’ anticipation, lanes start to form early. As Figure 4.12a illustrates, with group based perception, the agents at the edges starts forming lanes even earlier

than traditional RVO2 does. [128] also stated that two or three asymmetric lanes are more realistic than a larger number of symmetric lanes. The simulations seem to indicate that when IBP is used, the lanes produced by the RVO2 agent become more asymmetric and arguably more realistic.

Thus, it can be seen that, in this scenario, RVO2 with IBP qualitatively produces results that are at least as good as traditional RVO2. This is further confirmed by the average effort expended calculation shown in Table 4.4.2.3 where both models have similar values. The effort expended was calculated as in Section 4.4.1, using Equation 4.7.

4.4.2.2 Intersection (2 *mma*)

In this category, the typical scenario identified is an intersection of two corridors with groups of pedestrians crossing each other at the intersection. Figure 4.13 illustrates the results of the simulation. The figure indicates that group based perception does not make much difference in this scenario. Lanes are similarly formed in both scenarios with most agents passing through smoothly. However, due to the group based perception, the efficiency and symmetry of this movement is somewhat reduced with one agent being dragged out of its preferred path by the other group because it initially tries to avoid the group. However, the calculation of the average effort per person shown in Table 4.4.2.3 indicates that this is not that significant a difference. Moreover, it is difficult to argue which of the two results are more accurate without experimental data.

4.4.2.3 Circle (multiple *mma*)

The third and final category identified in [1] was also the most complicated in that there are several *mmas*. The scenario considered to best represent this category is one with 32 agents in a circular formation. Each agent tries to exchange positions with the agent diametrically opposite to it. Collision avoidance algorithms, like RVO2, are expected to smoothly navigate this situation without a deadlock developing. For IBP to be practical, it is expected that IBP does not adversely affect the effectiveness with which RVO2 handles this difficult scenario. As shown by the trails in Figure 4.14 and Table 4.4.2.3, this is indeed the case. Despite the trails produced being different in some ways, deadlocks do not appear and the average total effort per person is only marginally higher when IBP is used.

In this section, we have demonstrated the working of RVO2 using IBP by comparing it against the RVO2 with a traditional sensor range in three standard scenarios: bi-directional corridor, a two-corridor intersection and a circle scenario. IBP produced results that were at least as good as a traditional sensor range in all three scenarios. In the next section, we validate the proposed model against real world data.

Figure 4.13: Intersection Experiment: IBP seems to marginally reduce the efficiency with which collisions are avoided in this scenario. However, this is the case only for one agent. Moreover, it is difficult to determine which result is more accurate.

(a) RVO2+IBP

(b) RVO

Figure 4.14: Circle Experiment: IBP does not have a significant impact on the results produced by RVO2 in this scenario. Unnatural deadlocks do not develop in either scenario.

4.4.3 Experimental validation

The experiments so far discussed in this section, have been based on quantitatively and qualitatively analyzing simulations in different scenarios. However, to better validate the model and the theory of group collision avoidance on which it is based it would be necessary to compare it against real world data. Hu et al. [1, 129] conducted a series of controlled field experiments to study the interaction among individual participants

Model	Corridor (kJ/person)	Intersection (kJ/person)	Circle (kJ/person)
IBP+RVO2	152.725	141.183	115.802
RVO2	152.355	141.193	111.207

Table 4.2: This table shows the average energy expended per agent in standard scenarios and compares the results of RVO2 with traditional sensing against RVO2 with IBP. There seem to be no significant differences between the models by this measure.

Figure 4.15: A 13.5 m x 2.5 m area is demarcated for the experiments. The participants stand 10m apart and are directed to walk towards the other end at their preferred speed. (From [1])

and their movement in certain scenarios. In this section, the experimental setup used by Hu in his thesis is first reproduced. Following this, the results obtained are presented and compared against the results produced by RVO2 agents using Information Based Perception.

4.4.3.1 Experiment setup

This section explains the experimental setup and measurements made by Hu et al. to study human movement in particular scenarios⁴. In the original paper three experiments were conducted, since IBP does not produce results different from the underlying motion planning algorithm used (RVO2 in this case) for two of the scenarios only the first experiment analyzing movement of pedestrians walking towards each other is presented and considered in this analysis.

The experiments were conducted at an outdoor passage outside the School of Computer Engineering in Nanyang Technological University. Figure 4.15 shows the setup of the experiment. Each experiment consisted of three participants (marked 1, 2 and 3) moving at their preferred speed in the given direction until they reached the goal point marked by the border of the demarcated area. Hu et al ar-

⁴This section simply summarizes the experiments conducted by Hu et al. to provide a context for the subsequent analysis and validation done. For a more detailed discussion please refer to the original paper and thesis.

gue that this scenario is typical of several real life bi-directional movement scenarios. In order to shift attention from movement, the participants were given some simple algebraic calculations to do as they move. The experimental setup is shown in Figure 4.16. A video camera with adequate resolution was kept at fixed height above the experimental area using a horizontal tripod. The location of the camera and the lens was fixed to ensure maximum clarity and minimal distortion. Markers were kept on the experimental area to prevent participants from going outside the experimental area and to help in analysis of results obtained. Participants were also asked to wear white helmets so that the analysis from the video was accurate [130].

Figure 4.16: The experimental setup used by Hu et al. [1] for their controlled field experiments.

Videos of 20 experiments with 13 participants (9 male; 4 female) were

Figure 4.17: The data extraction process.

finally analyzed for characteristic movement patterns and other metrics. The process shown in Figure 4.17 was then used to extract the trajectories of each of the participants in each of the recorded videos. Briefly, the process was as follows: first the video was converted into a series of images at 20 frames per second which is an ideal frame rate to obtain a complete trajectory of the movement of participants. Once an image stack was obtained the image was scaled, rotated and transformed using the markers to remove distortions and bring it to metric scale. Following this, the next two steps quantize the image and remove noise such that, at the end of the process, each image only has black and white colors with the former depicting the background and the latter the participants. Lastly, an existing software is used for tracking each participant (white particle) in the quantized images. Fig. 4.18 gives an example of the whole process.

4.4.3.2 Quantitative measurements of experimental data

The data extraction process outlined in Section 4.4.3.1 produces for each participant in each experiment a trajectory. Trajectories that were similar were grouped together and normalized to find an average trajectory for each group. It was observed that these trajectories could be grouped into three main groups shown in Figure 4.19. The quantitative measurements that are explained next are all calculated on the average trajectory for each of these groups.

Figure 4.18: An example of the video extraction process

The different measurements used by Hu et al. are very neatly summarized in Figure 4.20. These metrics are:

- Trajectory length of participant i (L_i).
- Maximum deviation of participant i (D_i). This is the maximum value of dd_i^t for a participant over the course of the complete trajectory.
- The speed of the participant as function of time t (s_i^t).

Figure 4.19: The trajectories of participants can be grouped into 3 categories

Figure 4.20: Metrics proposed by Hu et al. for quantitative validation of models [1].

- The distance between agents i and j as a function of time t (d_{ij}^t).

4.4.3.3 Results and comparison

This section presents the simulation results produced using IBP when compared against the experimental data. The experimental setup was replicated by initializing the agents as shown in Fig. 4.21. Of the four metrics proposed, the maximum deviation and inter-agent distance against time seemed to differentiate best between the models. Hence, in this section, we use these metrics for evaluating IBP. Appendix A discusses briefly the results of the other two metrics.

Table 4.3 shows the maximum deviation values for the real world data.

Figure 4.21: The initial configuration of the simulation based on the experimental setup in real world experiments.

Maximum Deviation Calculation from Experiments			
	Category 1 (55%)	Category 2 (40%)	Category 3 (5%)
Agent 1	0.0785	0.320723	0.241166
Agent 2	0.8720	0.868398	0.320865
Agent 3	0.2656	0.18667	0.657496

Table 4.3: This table summarizes the maximum deviation according to experiments conducted by Hu [1]. Category 1 refers to the movement of Agent 2 to the left of the group, category 2 to the movement of Agent 2 to the right of the group and category 3 to the movement of agent 2 through the middle of the oncoming group.

It can be broadly seen that in 95% of the cases, Participant 2 avoided Participant 1 and 3 through either side. It can also be seen that the major deviation effort was taken by single agent ($\approx 0.87m$) rather than the group of two agents ($\approx 0.2 - 0.3m$). Table 4.4 compares the distance deviation calculation for the simulation results against the real world data. Due to the symmetry of the behavior of agents 1 and 3 in category 1 and 2 of the real world experiments, we do not lose much information by averaging category 1 and 2 to obtain the values in the second column of Table 4.4 which represents 95% of the real world data. The remaining columns in the same table show the results produced by traditional RVO2, Social Force and RVO2 with IBP. As can be seen, RVO2 with IBP produces the dynamics most similar to the real world experiments. This is more clearly illustrated in Figure 4.22.

Figure 4.23 illustrates the inter-agent distance as a function of time for the experimental data, RVO2, Social Force and RVO2 with IBP. As in the real world experiments of Hu it can be seen that agents 1 and 3 maintain a relatively constant value. It can also be seen that, at its closest point, as in the experiment, agent 2 becomes closer to one of agent 1 and 3 than they are to each other. This is closer to the experimental data than the results produced by RVO2 and Social Force using a normal sensor range as shown in Fig. 4.23d and Fig. 4.23e.

Thus, this section has demonstrated how Information Based Perception helps RVO2 produce results that are similar to those obtained by Hu et al. in their controlled experiments.

	Experiment (95%)	IBP-RVO2	RVO2	Social Force
Agent 1	0.199	0.0	0.043	0.372
Agent 2	0.870	0.855	0.003	0.0
Agent 3	0.226	0	0.038	0.372

Table 4.4: This table shows the trajectory length obtained for the simulation of the experimental scenario in [1]. The first column lists the average of category 1 and 2 which together comprised 95% of the real world experiment data of [1]. The same pattern of agent 2 having larger value than the other two is obtained only when IBP is used.

Figure 4.22: Graphical illustration of the absolute difference in maximum deviations that can be seen in Table 4.4.

4.5 Summary

In agent based simulations of crowds, agents go through a sense-think-act cycle. Traditionally, crowd simulations tend to focus on improving the thinking and acting parts of the cycle while using a simple spatial distance based perception system for sensing. As discussed in the literature and shown through the results in this chapter, in the context of movement simulation, this assumptions limits the ability of existing systems to produce realistic results.

Existing theories of human cognition were used to develop an Information Based Perception model for agents which is based on perceived information rather than spatial distances. Through simulations and comparison against real world experiments, it was shown that this is a more appropriate model of human perception for crowd and egress simulation as it helped produce more realistic group avoidance behavior. The idea that humans have limited perception capacity such that they only process certain obstacles more relevant to collision avoidance was incorporated and its usefulness demonstrated. The real world experiments conducted by Hu et al. were presented; comparison of simulation results against experimental results showed that Information Based Perception can help models like RVO2 replicate the results produced by the experiments.

The IBP model can be extended so that during filtering it can recognize cues that contain event and environment information as well. Modeling cues and their impact on the early parts of the evacuation process is a crucial problem that has been ignored by crowd simulations to date. The idea of modelling different cues and their effect on the agent's information processing capabilities as suggested by Kuligowski [35] is investigated in more detail in the next chapter.

Figure 4.23: Plot of inter agent distances over time for various models and the experimental data.

Chapter 5

Modelling Pre-evacuation Behavior

Ideally, when a fire starts a fire alarm goes off; all occupants hear this alarm and use the nearest safe exit to leave the building. However, as discussed in Section 2.2.1 and demonstrated tragically at Seton Hall University [28] this hardly ever happens. In many cases, occupants are desensitized from hearing false alarms and often do not start to evacuate until they are completely sure that egress is necessary. This uncertainty about the authenticity of the first sign of danger isn't an isolated incident [24, 43, 44, 46–50]. This delay in reaction time of evacuees can have a huge impact on existing plans for egress if it is not taken into consideration. Thus, it is important to consider this when developing evacuation plans. In other words, when studying the behavior of evacuees, it is necessary to study and understand their actions from the time at which the fire started right up until the point where the last person evacuated [47].

As discussed in Chapters 2 and 3, pre-evacuation uncertainty and investigation are features of human behavior during egress that are rarely considered in existing models. To recap, *pre-evacuation* refers to the period of time that elapses after the start of the fire alarm before the person starts evacuating. While some models [29] do have a simplified model of pre-evacuation behavior, they fail to model it in enough detail to enable their extension to more general cases. For example, a fire alarm could have different effects based on the clarity and believability of the alarm [17, 21]. This variability is hard to simulate in existing models of pre-evacuation behavior. Also, during an evacuation people exchange event and environment related information with others. Evacuees are unlikely to follow every message that they receive blindly. There is a variability in the *trust* in messages received that can have different effects on egress. This is rarely considered in existing models.

In this chapter, we present a model for simulating pre-evacuation behavior and event identification¹. In the model, the evacuees identify and process information in terms of event cues which exist throughout the environment. Section 5.1 first explores some of the existing work that motivates the need to model

¹This model was presented at the Pedestrian and Evacuation Dynamics Conference in 2012 and appeared in print in 2014 [107].

pre-evacuation behavior. Following this, Section 5.2 explains the proposed model and, finally, Section 5.3 illustrates through experimentation the usefulness of this model, the importance of modelling pre-evacuation behavior and having a communication model.

5.1 Related Work

As discussed in Section 2.2, there are a lot of conflicting theories on how humans behave in emergencies and why they behave as they do. However, there are certain behaviors that are fundamental to human nature and are generally accepted as fact, such as the constant search for information [22, 26, 46, 47]. This section first summarizes the existing knowledge of human behavior during egress with special emphasis on pre-evacuation behavior. Following this, some existing models of pre-evacuation behavior and evacuee communication are presented.

5.1.1 On modelling pre-evacuation behavior

Several studies of human behavior during emergency egress [26, 51, 68], have shown that an evacuee's first reaction after realizing that there is an unusual situation is to investigate and gather more information about the situation. Evacuation starts only once the need for evacuation is established. *Cues* are the key to understanding this transition from realization to investigation and, eventually, to evacuation.

Cues are certain changes in the environment that indicate that something is wrong or different from normal [22]. They come in different forms. Fire and smoke are the typical and most effective cues for an evacuation. Fire alarms and people running or instructing to escape are also cues. According to Sime [22], there are three kinds of cues:

- Ambiguous Cues: E.g. Hearing noises or shouting, or seeing someone run.
- Verbal Cues: E.g. Instructions from a companion, announcement from the stage.
- Unambiguous Cues: E.g. Seeing smoke or fire, or seeing someone run with a fire extinguisher.

According to some researchers [44, 51], an ambiguous cue by itself does not cause a person to initiate investigation. Rather, the cue has to persist for a period of time before investigation begins and results in the finding of an unambiguous cue. Interestingly, according to Tong and Canter [47], even unambiguous cues don't result in people immediately exiting the building, rather it initiates a complicated process consisting of information searching and affiliation.

The effect of cues and the evacuee's reaction to it depends on many factors. The identity of an individual's primary group and its proximity and availability determines the reaction of a person to a cue [22]. Various studies [21, 37, 43, 46, 47, 53] emphasize the importance of the location and the person's role in the significance of cues. As an example, a fire alarm at home is more likely to cause a person to act than a fire alarm at their office, which will most likely be considered a false alarm. People's societal role determines their training and responsibility and thus their alertness to cues and the preparedness for reacting to it.

Their groups, location and environment aren't the only factors that influence people's behavior during evacuation. There are also a lot of *intrinsic factors* that influence how people react to fires. What these factors are and how influential they are are a matter of much debate. Over the years there have been several surveys [35, 37, 47] that discuss these intrinsic factors.

Andrée and Eriksson's report [36] even had a cross cultural study that compared the evacuation behavior of Swedish students against the behavior of Australian students. Except for grouping behavior they found few significant differences in the behavior during evacuations. Kobes et al., in their survey [17], compared some studies of evacuation behavior from the USA, Great Britain and Australia and commented on them being "identical in their essence". Some factors like age and gender were found to not significantly affect pre-evacuation behavior. A person's social role is one of the commonly accepted factors that influence evacuation behavior [17, 21, 37]. Factors like the person's experience with fires, training, disabilities, familiarity with the environment, etc. are accepted to have a great influence on pre-evacuation behavior.

Close to a hundred different studies on human behavior in fires were used by Kuligowski [35] to compile a list of the factors that influence egress and more specifically pre-evacuation behavior. In this article, she suggests that the period that we term as *pre-evacuation* itself consists of two phases. Phase 1 is called *perception*; The idea being that the existence of a cue does not guarantee its perception. The Table 5.1 lists some factors that can affect this perception of a person and whether these factors increase or decrease the chance of the person perceiving the abnormality in the situation.

Kuligowski calls the next phase *interpretation*. During this phase, the person searches for more information to verify whether a fire has actually started and if it actually poses a threat that needs to be handled. Many studies [26, 51] have confirmed the importance of this phase, though sometimes they are known by different names. In [26] this phase is called *unconflicted inertia* to indicate how the person either continues what he's doing or tries to finish of his activity without actually beginning to evacuate. In [47], this phase is called *investigation* to indicate the search for information that enables the interpretation of the system. Regardless of what

it is called, this phase consists of two parts: 1. defining the situation as a fire and 2. defining the risk that the situation poses.

Kuligowski categorized the factors that influence these phases into two types: occupant based factors (which are equivalent to the intrinsic factors mentioned earlier) and cue based factors. These factors and their effects are shown in Table 5.1. Increases/Decreases signifies a cues effect on that particular phase. Also, cue based factors, do not simply consist of events, but also their nature and features.

One of the factors that encourage the programmatic implementation of cue based factors is the fact that the effect of a cue can be explained to be caused by the nature and characteristics of the cue rather than the specific cue. In other words, each cue can be described in terms of its ambiguity, consistency with other cues and its source and it is this description that determines the effect of the cue.

5.1.2 Existing models

There are very few existing models that take pre-evacuation period into consideration. Pre-evacuation decision making of an individual was modeled by Pires [31] using a simple Bayesian Belief Network (BBN).

França et al. [41] created a simulation model of the development of panic behavior during emergency egress. This model implemented the hysterical belief theory [20] and modeled how panic first develops and then evacuation happens. It also had a basic communication system through which agents exchanged mood information (which is a key factor in the development of panic) by using the grid based environment as a medium for communicating messages. Despite pre-evacuation behavior being modeled in some detail, it is not possible to extend this model to replicate the heterogeneity in people's reaction to cues.

ESCAPES [29] is a model that takes into account some factors like the spread of knowledge, fear and emotion between the different evacuees. These factors are used to create a simplistic model of pre-evacuation behavior. Effectively though, only the *perception* part of pre-evacuation behavior is modeled. This means that once an agent perceives a cue i.e. another agent running, the perceiving agent immediately starts to evacuate.

The event identification and communication model proposed in this chapter have been influenced by these models. However, it is unique in the way that the diversity of cues and their effects can be considered; and more importantly, in how both the perception and investigation phase of pre-evacuation behavior can be simulated in this computational model. Data on pre-evacuation behavior is difficult to gather: If participants are given prior warning like in fire drills, their behavior is likely to be unrealistic; on the other hand, not giving prior warning could be risky and dangerous. Thus, in this chapter, the proposed models are validated against established theories.

Table 5.1: The list of factors affecting evacuation behavior and their influences as presented in Kuligowski’s survey [35]

Factors	Perception	2a: Definition of the Situation as a Fire	2b: Definition of the Risk to Self/Others
Occupant-based pre-event factors			
Has experience with fires	Increases	Increases	Increases
Has knowledge of fire/ training	Increases	Increases	Increases
Habituation with environment	Decreases	— ¹	—
Has knowledge of routes	—	—	Decreases
Has frequent experience with false alarms	—	Decreases	—
Has a feeling of security in the building	—	Decreases	—
Has perceptual disability	Decreases	—	—
Is older adult	Decreases	—	Increases
Is woman	Increases	—	Increases
Speaks the same language as others	Increases	—	—
Has frequent interaction with family	Increases	—	—
Occupant-based event factors			
Has a higher stress/ anxiety level	Decreases	—	—
Perceives a time pressure	Decreases	Decreases	Increases
Presence of others (especially loved ones)	Decreases	—	Increases
Proximity to fire / Visual Access	Increases	—	—
Sleeping	Decreases	—	—
A higher number of behavioral processes(>1)	—	Increases	—
Defines situation as a fire	—	N/A	—
Cue-based factors			
A higher number of cues	Mixed ²	Increases	Increases
Consistent cues	—	Increases	Increases
Unambiguous cues	—	Increases	—
Social cues (others’ actions) that are consistent with an understanding of a fire situation	—	Increases	Increases
Official source	Increases	Increases	—
Familiar source	—	Increases	—
A higher dose of toxic gases	—	Decreases	—
Extreme/ dense cues	Decreases	—	Increases
Visual/ audible cues	Increases	—	—
Risk Information	—	Increases	—

¹ Areas where no research is found is marked by —² Research conflicted on the direction of influence of this factor

5.2 A Bucket Model of Event Identification

Section 5.1.1 discussed how ambiguity, source and consistency of the cue [22, 47, 68] are the key factors in determining the effect of a cue. This is the fundamental idea used in the proposed model. Ambiguity refers to the clarity of the cue. A cue that gives a clear signal of an event would be unambiguous while an ambiguous cue will lack this clarity. For example, a traditional ringing fire alarm is highly ambiguous while the modern fire alarms with voice updates are much less ambiguous and they have different effects on pre-evacuation behavior. The source of a cue is also important; there is a clear difference in perceived danger in a situation when the warning is given by figure of authority like a fire marshal as opposed to an unknown stranger. Lastly, if the signals given by cues are inconsistent then it would take longer for a person to start egress. For example, if no other occupant is perceived to be raising an alarm or trying to evacuate, most people would not react to smoke [52].

Each cue is located at a particular location (x, y) in the environment and is sensed by agents when within their perception range. When a cue is sensed, an agent obtains a certain amount of information μ about the current state of the world. Thus, a cue in the Bucket model is represented as a tuple (μ, x, y) where μ is the amount of information obtained from the cue and (x, y) is the location of the cue. When an event or an action of significance occurs at some location a cue is created at that location.

Once perceived, these cues are processed by the agent's Event Knowledge System which manages the information from cues. The Event Knowledge System has two *buckets* of information: one corresponding to *investigation* (β_i) and the second corresponding to *egress* (β_e). When a cue is perceived, a certain amount of information (μ_{β_i} and μ_{β_e}) is added to the respective bucket(s) based on the ambiguity level, consistency and source of the cue. For each bucket, a *threshold* (τ_{β_i} and τ_{β_e}) is initially fixed based on the intrinsic characteristics of the agent. When the amount of information in a bucket overflows the threshold the agent changes its internal state and proceeds towards a new goal based on the behavior model that is used. For the purpose of the experiments in this chapter, each agent has three states that it can be in: *normal* behavior where the goal is the center of the *starting* room of the agent; *milling* behavior where the agent gathers with other agents at the nearest *corridor* (See Figure 5.2); and *evacuation* where the agent heads towards the exit through the shortest route. Figure 5.1 illustrates the working of the bucket model within this simple behavior model.

There are two sets of values here which have been mentioned: bucket thresholds (τ) and the amount of information (μ) that a cue contributes. As mentioned earlier, each cue is characterized by its consistency, ambiguity and source. Cues that indicate a need to egress have $\mu > 0$ and add information to buckets while cues that do not, remove information from it (i.e. $\mu < 0$). Thus if an inconsistent set of cues

Figure 5.1: An illustration of the bucket model and the behavior model of agents used in the experiments. Each cue adds (or removes) a certain amount of information μ to one or both buckets (β_i and β_e). An overfilling bucket causes a state transition.

are observed at a given time, less or maybe even no information might be added to each bucket.

The amount of information that is added or removed (i.e. the value of μ) depends on the ambiguity and source of the cue. The key value in the *bucket model* is the ambiguity of the cue. Each cue has an ambiguity value that is represented as a real number in the range $[0, 1]$. Since each cue adds or removes information from the *investigation* and *egress* buckets, there are two information values associated with each cue, one for each bucket.

In theory, the amount of information added by a cue can be any decreasing function of the ambiguity of the cue. For simplicity, we assume that it is a linear function with the value added by a cue of ambiguity α to the bucket β being given by:

$$\mu_\beta(\alpha) = \begin{cases} \tau_\beta \left[1 - \frac{(\alpha_\beta^{max} - \alpha)}{\alpha_\beta^{max}} \right]^{-1} & \text{if } \alpha \leq \alpha_\beta^{max} \\ 0 & \text{if } \alpha > \alpha_\beta^{max} \end{cases} \quad \text{where } \beta \in \{\beta_i, \beta_e\} \quad (5.1)$$

Where τ is a constant equal to the egress bucket threshold value for a normal agent and is determined by the maximum time for the bucket to fill and the intrinsic characteristics of the agent. A higher value for τ would mean more information is added per cue and the bucket would fill faster. In terms of the formula, it simply determines the slope of the curve and the range of μ . We arbitrarily choose a value 10,000 for τ so that the rate of change of μ for every 0.1 unit change in ambiguity α is greater than 1.

α_e^{max} refers to the maximum ambiguity level that has effect on the egress bucket which in this case is fixed at 0.9 since we assume that cues more ambiguous

than this do not add any information. This helps us simulate those cues that only result in investigation even if they persist for a long time and never trigger evacuation itself.

A similar procedure is used to determine μ_i . An ambiguity value of 1.0 in this case represents a small amount of information, as opposed to no information in the case of the egress bucket. This is because even an unambiguous cue when perceived for long enough initiates investigation [47]. Thus, $\alpha_i^{max} = 1.0$.

The value for threshold of each bucket is equal to τ for normal agents. For special agents, and to consider certain special scenarios and effects, we can simply choose a threshold $\tau^s < \tau$; this will be demonstrated in the experiments in Sections 5.3.2, 5.3.3 and 5.3.4.

The effect of the source of the cue is seen in the communication model. Communication is implemented as *messages* sent from one agent to the other. Each message has a message cue and environment information. The message cue works just as other cues with it's own ambiguity level. Here the term ambiguity is used to refer to the trustworthiness of the source of the information. In the current version of the model, we use a simple communication mechanism whereby agents only exchange knowledge about the state of the fire, or more precisely the inaccessible parts of the map (those blocked by fire).

In the experiments in this chapter, a simple cellular automata based fire model and a simple finite difference smoke model are used to simulate the propagation of fire and smoke². This creates fire and smoke cues at locations near the fire which spread as the simulation proceeds. The environment itself is made up of various rooms which are connected by links. Some rooms that link multiple rooms are marked as *corridors* and are the gathering points for the agents when investigating. All agents are assumed to have complete spatial knowledge of the environment and are initially located in their *starting* rooms. As soon as the fire starts, fire alarm cues are placed all over the environment. Agents react to these cues and mark observable pathways that are blocked as inaccessible in their spatial knowledge graph. Agents die immediately on contact with fire. Smoke, depending on the concentration, initially slows and eventually kills agents. All agents either escape or are killed at the end of the simulation.

5.3 Results

This section discusses the experiments that were conducted to demonstrate the effect that cue perception, pre-evacuation behavior and communication can have on egress. The simulations were conducted on the two floor office environment shown in Figure 5.2. This was adapted from the first two floors of the World Trade Center in

²Validated models of fire and smoke are not necessary for testing the proposed model- a plausible model is sufficient

(a) Ground Floor

(b) Second Floor

Figure 5.2: The two floors used for the simulation which is a modified version of the first two floors of the World Trade Center in California. The corridors, staircases and exits on both floors are marked. The numbers are the room numbers which are used in the experiments.

California. Simulations were conducted with 200 agents randomly distributed over the two floors of the building and data was collected after averaging over 20 replications of the simulation.

As shown in Figure 5.2, the layout of the building is complex. As a result, different evacuation dynamics may be produced when the fire is started in different rooms. However, with close to a 100 rooms, running simulations with the fire starting in each room would be time consuming and, more importantly, unnecessary. Rooms close to each other or at a similar distance from the key locations (staircases and the exit) tend to produce similar dynamics. So we first grouped the rooms based on the evacuation dynamics produced in order to identify the experiments that need to be conducted. We did this by evaluating survival rates (percentage of occupants who survived) in scenarios of two types.

In the first type of scenario, we assume that all agents start evacuating immediately on hearing the fire alarm (i.e. at the start of the simulation). This is the best case scenario. Pre-evacuation behavior will increase total evacuation time and reduce survival rate. As with the other experiments, there are 200 agents randomly distributed over the two floors. Ten replications of the simulations were run for the fire starting in each room and the simulation was completed when all agents escaped or died. The survival rate was determined for each run. Figure 5.3 shows the results of this experiment as a heat map. The number on each room is the survival rate for that room. The standard deviation for all rooms were between 0 – 3 percentage points and is not shown in the figure. As expected, if the rooms are close to the exit or the staircase closer to the exit, then the survival rate is generally quite low. As the starting room moves further from the exit there is a higher chance of survival.

There are some rooms close to the exit with lower than 40% survival rate despite immediate evacuation. These can be ignored in our analysis, since modelling pre-evacuation behavior would not make any difference to the survival rate if the fire is started in these rooms. There are probably other rooms which can also be ignored due to the lack of effect of pre-evacuation behavior. In order to determine such groups of rooms in a more systematic manner, a second set of simulations were run with the agents only evacuating when they actually perceive the smoke (or fire). The results of this are shown in Figure 5.4. The pattern is mostly similar to the Figure 5.3 with just much lower survival rates and most survival rates being below 50% since the occupants of a different floor generally perceive the fire/smoke too late.

To understand the pattern of differences between Figure 5.3 and Figure 5.4 more clearly, the difference in survival rates in the two scenarios was found and plotted in Figure 5.5. The survival rates in this figure, along with the location of the room with respect to the two staircases and the exit to divide the rooms into 10 groups as shown in Figure 5.5. The rooms colored black i.e. the corridors and the ones with very small differences were ignored and not used for analysis; this also includes the rooms on the bottom left corner of the second floor which was so far from the exit that most occupants inevitably survived. For each of the other groups a stereotypical room was selected which had a survival rate somewhat close the mean

(a) Ground Floor

(b) Second Floor

Figure 5.4: This figure illustrates the survival rate when the fire is started in different rooms and an agent starts evacuating only when it perceives the fire or smoke. The number on each room is the survival rate for that room.

(a) Ground Floor

(b) Second Floor

Figure 5.5: This figure shows the classification of rooms into groups based on the difference in the survival rates shown in Figure 5.3 and Figure 5.4. The color of the groups are superimposed on the heat map. Corridors and rooms with very low difference are not considered to be part of any group since they are not considered in the experiments.

of the group and which was located somewhat centrally within the group. The rooms chosen for further analysis were: 14, 30, 53, 71, 81, 190, 224, 240, 252 and 377. Each of the following experiments were conducted with the fire starting in each of these rooms to understand the impact of room location on the results.

5.3.1 Experiment 1: the effect of fire alarm ambiguity

A fire alarm with a simple ringing sound is much less clear and more ambiguous than a public announcement system that explicitly states that it is not a drill and gives real time updates about the situation. Thus an unambiguous fire alarm will ensure a faster evacuation and a higher survival rate. In the first experiment, we demonstrate the effectiveness of the proposed model in modelling this scenario. We make the reasonable assumption that the fire alarm can be heard clearly at every location on the map and place cues in every room.

To examine the effect of this difference in clarity of fire alarms, the experiment was repeated for different values of ambiguity (from 0.0 - 1.0). In this simulation, message cues are not simulated and the only other cues other than the *fire alarm cues* are the *fire* and *smoke Cues*. A fire cue is a clear indication that there is a fire and an urgent need to start evacuation, so it is given an ambiguity of 0. Since a smoke cue is not as clear as a fire cue we arbitrarily assign an ambiguity of 0.2 to it.

Figure 5.6 shows the percentage of agents that survived as a function of fire alarm ambiguity. As expected, regardless of the room in which the fire is started, there is a drop in number of survivors as the ambiguity of the alarm increases. This is because a less ambiguous alarm, gives more information, results in the buckets overflowing quicker, a shorter pre-evacuation period and obviously faster evacuation. A more ambiguous alarm will have to be perceived for longer before it fills the bucket and triggers a state change. This delay will result in occupants further away from the exit getting insufficient time to escape depending on the room in which the fire is started. As observed earlier, the further the fire is from the exit and the staircase near the exit, the higher the survival rate.

Certain cues never trigger evacuation and, at best, trigger investigation if persisting over a period of time. In order to model this phenomenon, we explained in Section 5.2, how α_e^{max} is fixed at 0.9. This means that for cues with ambiguity $\alpha > 0.9$, no information is added to the *egress bucket* β_e . This results in the sharp drop at $\alpha = 1.0$ in Figure 5.6 as agents only start evacuating when they actually observe the fire or smoke cues themselves. Even in this scenario, however, the survival rate is higher than shown in Figure 5.4 because milling behavior still takes place. The effect of milling is analyzed in more detail in Section 5.3.3.

The flat curve for room number 252 is interesting but not surprising. This is because of the peculiar location of the room far from both staircases and on the second floor. Thus, for any value of $\alpha < 0.9$, most evacuees have enough

Figure 5.6: This figure shows the percentage of evacuees that survived as a function of the ambiguity with the fire being started in different rooms. While there are obvious inter room differences, it can be clearly seen that the more ambiguous an alarm, the lower the survival rate.

time to escape before the fire and smoke blocks the crucial staircases and exits. For $\alpha > 0.9$ however, as explained previously, egress does not start until the smoke or fire is directly perceived and results in a sharp drop in survival rate.

This experiment has highlighted the importance of modelling and considering pre-evacuation behavior and the value of having a clear and unambiguous fire alarm. While, the effect of ambiguity of fire alarm and pre-evacuation is clearly seen independent of the location of the fire, the location does affect the extent of this effect.

5.3.2 Experiment 2: informed and informing agents

In a situation where it is possible to inform a certain number of occupants of a building about the urgent need to evacuate and these occupants inform others as they evacuate, an interesting question might be to examine how many people need to be informed for this strategy to be effective. To examine this, we divide the set of agents into two groups: normal agents and informed agents. The former behaves as in the previous; the latter start evacuating as soon as the simulation starts and broadcast messages to other agents to start evacuating immediately. This immediate evacuation can be modeled by setting the threshold of both the investigation and the

egress bucket to zero thus ensuring that the phase is triggered right after the first time-step of the simulation. Message cues that have an informed agent as their source have an ambiguity of 0 indicating that these messages are completely trusted by other agents. The fire alarm ambiguity cue is set to 1.0 to minimize its effect.

Simulations were conducted with the proportion of informed agents varied from 10% to 100% of the 200 agents. Experiments were run with the fire starting in each of the stereotypical rooms. The blue curve in Figures 5.7a, 5.7b, 5.7c, 5.7d, 5.7e, 5.7f, 5.8a, 5.8b, 5.8c and 5.8d shows the results of this experiment.

Except in the case of the fire starting in room 377, with just 15% agents being informed the survival rate increased two fold or three fold in almost all cases. Room 377 being so close the vital staircase that is close to the exit, has less than 50% survival regardless of the situation (As also seen in Figure 5.3 and Figure 5.6) and expectedly shows a much slower rate of improvement than the other rooms. Also except in the cases where the fire starts in room 377 and 53, survival rate is greater than 50% and, in most cases, greater than 60% with just 20% of the occupants being informed.

In order to isolate the effect of these informed agents propagating information, the red curve in the Figures 5.7 and 5.8 shows the effect of these agents only being *informed* agents that evacuate immediately without *informing* other agents. The importance of these evacuating agents broadcasting the information they know is highlighted by the 20-30% difference in survival rates that are observed when 20-60% of the occupants are informed. This difference is lesser but still significant even in the case of room 377, 53 and 71. What is common to these rooms are their proximity to the staircase closer to the exit. This is probably because when this staircase is blocked the shortest route towards the exit becomes longer. This results in more deaths.

This section demonstrated the usefulness of informing and possibly training a certain proportion of the agents about the need to evacuate. Even when randomly distributed in the environment, less than a quarter of the occupants of a building need to be trained for significant improvements to be observed. More strategic placement of the informed occupants would obviously mean fewer of them are required. The usefulness of strategic placement of informed occupants is considered in more detail in Section 5.3.4.

The experiments also revealed the importance of these informed occupants broadcasting their knowledge as they escape. Again, the scenario tested in this experiment was quite simple in the sense that these informed agents informed only those occupants that it met while moving towards the exit. Their effectiveness will probably be much higher if they actively tried to inform other agents. However, in that situation, there is the added consideration of how much effort should they make in finding and informing others before making their own way to the exit. However, one of the reasons for the effectiveness of these informed agents is that most of the

Figure 5.7: These figures show the effect of informing a certain set of agents to evacuate immediately. The blue curve shows the survival rate when these informed agents broadcast and inform other agents as they evacuate. The other curve helps contrast this with the situation where the informed agents simply evacuate without informing other agents, thus illustrating the importance that informed agents also inform others as the evacuate.

uninformed occupants would have been *milling* in the corridors and were thus easily reachable; this probably alleviates the need for the informing agents to actively try and find and inform other agents. In the next experiments, we examine the importance of this milling behavior.

Figure 5.8: These figures show the effect of informing a certain set of agents to evacuate immediately. The blue curve shows the survival rate when these informed agents broadcast and inform other agents as they evacuate. The other curve helps contrast this with the situation where the informed agents simply evacuate without informing other agents, thus illustrating the importance that informed agents also inform others as the evacuate. (contd.)

5.3.3 Experiment 3: is milling important?

The milling effect of people gathering in a common location to investigate and confirm the need to evacuate has been reported by interviewers of victims [50, 60] and is also one of the main features of the evacuation behavior modeled in this chapter. However, it is also a feature that is generally ignored in most models. Even in the rare cases that do consider pre-evacuation behavior, it is simply modeled as delay in evacuation. In this section, the importance of modelling milling behavior is demonstrated.

For this experiment, we compare agents that mill against agents that don't mill and simply move straight to evacuation. We model this by creating a new kind of agent that has an investigation threshold $\tau_{beta_i} = +\infty$ is modeled. These agents will have their egress bucket β_e overflowing before their investigation bucket β_i . Thus, they transition straight from normal behavior to evacuation (Figure 5.1). We compare these agents against the original agents to test the hypothesis that milling

Figure 5.9: The figure illustrates the difference created by and the usefulness of milling behavior by showing the increase in survival rate that is observed when there is milling against when there isn't milling/ pre-evacuation behavior.

behavior has a significant effect on evacuation characteristics. Since both agents have the same threshold for the egress bucket they will both start evacuating at the same time and the normal evacuation characteristics are identical.

Milling is important and useful in situations where evacuees communicate with each other. So, in the first experiment we fix the message ambiguity as 0.0 and fix fire alarm ambiguity at 1.0 so that agents only start evacuating, either if they see the actual fire or smoke or if they are told about the fire by evacuating agents. We only do the analysis on agents spawned in the first floor because the agents in the second floor have no way of knowing about the need for evacuation until it's too late. Figure 5.9 compares the survival rates with milling against the situation where there is only delayed evacuation. The figure clearly shows milling improves survival rate in all but the three cases where the fire is started in rooms 224, 377 and 53. In these situations, most agents that do manage to escape, take the route near the top-right staircase. In this route, they do not meet other occupants who get trapped without knowing about the fire (or finding out about it too late). Thus milling makes no difference in these three scenarios.

5.3.4 Experiment 4: the importance of spatial knowledge propagation

In experiment 2, we considered the effect of having a certain proportion of the occupants as trained personnel who are informed immediately about the need to evacuate. A simplistic scenario where these informed agents are randomly distributed over the environment was modeled and their effectiveness demonstrated. Significant improvement was seen because the evacuating informed agents not only informed other agents about the need to evacuate but also transmitted information about the areas that they knew were inaccessible. This is a significant advantage that a set of trained/ informed personnel can have over the most expensive or unambiguous fire alarm. Unambiguous fire alarms like loudspeakers or PA systems constantly giving instructions would still find it very difficult to transmit spatial information. Thus despite getting everyone informed and evacuating, no one would know where the fire is. In this experiment, we test whether spatial knowledge propagation through informed personnel can be an improvement over an unambiguous fire alarm.

The informed agents we consider in this experiment are different from those in Experiment 2 in two ways. Firstly, informed agents are modeled such that they can communicate spatial and cue information to other informed agents immediately. This is a reasonable assumption since it is likely that trained personnel will have some method for communicating efficiently with each other. Secondly, rather than being randomly distributed over the environment, they are located such that each corridor has one informed agent. This strategic distribution increases the likelihood that at least one informed agent perceives the fire and knows its location and propagates this information across the informed agent network.

The usefulness of spatial knowledge propagation would be best seen in a situation where survival rate is low due lack of spatial knowledge. Thus, we chose room 53 as starting point of the fire based on the results of the fire alarm ambiguity experiment (Figure 5.6). An empirical observation of the simulation of this scenario confirmed that one of the major reasons for dropping survival rate was the lack of spatial knowledge. Evacuees would realize too late that the staircase close the exit was blocked and would not have enough time to reroute and escape. Figure 5.10 illustrates the result of running 20 replications of the scenario with strategically distributed informed agents and contrasts this with a situation without these agents.

The simulation results seem to indicate that in scenarios like this, informed agents do indeed make a difference. Close to 60% survival rate is observed in scenarios with informed agents regardless of the ambiguity of the alarm while normally survival rate would fall to below 20%. More interestingly, there seems to be some improvement in survival rate even when compared against an unambiguous fire alarm. In fact, even in situations with a highly ambiguous fire alarm, strategically placed informed agents seem to have as much effectiveness as an unambiguous fire

Figure 5.10: Simulation results indicate that strategically placed informed agents can provide a significant improvement in evacuation efficiency and is as effective as an unambiguous fire alarm in ensuring higher survival rate.

alarm would have.

5.4 Summary

A review of existing literature and incidents, at the beginning of this chapter, revealed the crucial role played by pre-evacuation behavior in determining egress dynamics and survival rate. Subsequently, a survey of existing computational models revealed that pre-evacuation behavior is never considered in any detail in existing computational models. Thus, a novel cue modelling and perception system which enables the detailed modelling of pre-evacuation behavior was developed.

This model is based on the work of Kuligowski [35] and is based on the idea of perceived cues adding information to two information buckets whose overflowing results in phase changes. Occupants first start an investigation phase and once the need for evacuation is confirmed, they start evacuating.

It was demonstrated through simulation that modelling pre-evacuation behavior and the effect of ambiguity of perceived cues can have significant effects on egress dynamics and efficiency. The value of training a certain proportion of the

population and the critical number of these trained personnel needed was identified. Simulations indicate that strategic placement and training of a small subset of occupants might be cheaper but just as effective as an expensive and unambiguous fire alarm. The value of milling behavior and the need for modelling this was also demonstrated through simulations.

The proposed pre-evacuation behavior model can be a valuable tool for studying and analyzing many more real world issues. Similar simulations can be used to do safety by design analysis of buildings by ensuring that milling/gathering points are efficiently located. Other safety-by-design elements might have to do with budget constraints during large events or festivals. Strategic placement of more expensive but less ambiguous fire alarms along with some cheaper more ambiguous ones can reduce the cost while ensuring the effectiveness of emergency systems during these events.

While the model is based on theories that are based on studies of real world incidents, the results produced have not been validated against real world data. In the absence of detailed video footage or at least statistics of people evacuating from a building, validating a model of pre-evacuation behavior is very difficult. This is especially so, because in fire-drills and practices the observed pre-evacuation dynamics will be very different because the evacuees generally have prior knowledge of the drill.

Setting up a controlled experiment is difficult due to the dangers involved to participants. A possible way to tackle this difficulty in getting real world data would be to make use of virtual environments like games. This could be used to improve our understanding of cues and their effect on pre-evacuation behavior in general. We explore this approach further in the next chapter where a game based approach is used for studying human exploration behavior. In the simulations in this chapter it was assumed that each agent had complete knowledge of the environment, at least in terms of every room and how to get to each room. This is almost never the situation in the real world and to improve the efficacy of simulations it is necessary to model this partial knowledge. However, our current understanding of how people explore a partially known environment and how their knowledge develops over time are both very limited. The next chapter presents a game based methodology for studying this.

Chapter 6

A Game-based Investigation into the Role of Memory in Human Exploration and Navigation

Several studies have shown that during egress people tend to take routes that are familiar to them rather than the shortest route [21, 37, 44, 45]. This is because, in most cases, occupants of a building do not have a complete cognitive map of its layout. This is sometimes the case even after the occupants work in the building for several months [131]. However, existing agent based models of egress, generally make the simplifying assumption that occupants have complete knowledge of the environment. In the rare cases like [30] where partial knowledge is modeled, it is simply assumed that people have an eidetic memory and never forget a room that they perceive once. Moreover, it is arbitrarily assumed that agents explore unknown environments in either a depth first or a breadth first manner. How people explore unknown environments and gain spatial knowledge from this will obviously have an impact on the paths that are taken during egress. However, a major difficulty in modelling spatial knowledge accurately is our limited understanding of the role of memory in indoor wayfinding and exploration. In this chapter, we try to address this issue.

How humans gain and store knowledge of their surroundings has been an area of research for several decades now [132–136]. A significant amount of research has been conducted in trying to understand how wayfinding is done using this knowledge in outdoor environments [134, 137]. However, there has been much less work done on indoor exploration [136, 138–140]. With the increasing number of shopping malls, airports, high rise office buildings and residential towers, there is an increasing likelihood that the occupants of a building may not be regular visitors and as a result will have very little knowledge of the layout. If in such a situation, there is a need to evacuate the building, it is important for planners to know how the occupants would react and find their way to the emergency exits. However, there are several aspects

of how people explore an unfamiliar environment and the interaction between exploration and memory during this process that is still not understood. More importantly, we believe that existing methods for studying this interaction are limited in certain ways.

Understanding how people navigate and explore unknown spaces is scientifically challenging. One of the primary reasons is that experimentally studying such processes is difficult and time consuming. Even after an experiment is designed, it is often difficult to scale the experiments to many participants and therefore account for various forms of sampling bias. This can be seen in several existing studies [136, 139, 140] which, despite having experiments that are excellently designed, suffer from having few participants and thus probably being susceptible to sampling bias.

Recently, Bode and Codling [141] demonstrated a possible solution to this problem by using a point and click game to study exit route choice during evacuation and getting more than 150 participants to participate in their experiment. Recent years have seen the emergence of serious games (or gamification) as an approach for conducting human subject experiments [39, 142–144]. This approach can elicit the same behavioral response as real world experiments [145] while reducing the overhead associated with conducting large-scale real world experiments. In the past this type of approach has been used for validating crowd simulation models [108, 142, 146] and understanding behavioral responses to dynamic information during egress [11, 147]. Serious games [143] have a long history in training, for example, aircraft simulation [148] and medical training [149].

Due to the nature of the data under consideration, namely spatio-temporal tracks of many individuals, *desktop virtual reality experiments* are the most practical way to conduct such experiments [139, 140]. For exploration in particular, the use of a virtual environment allows the experimenter to design particular structures and understand how they affect people’s decisions. We believe that delivering the experiment through a game, rather than as a standard social experiment, is an ideal way to scale to large sample sizes.

In this chapter we adopt the same methodology to gain an understanding of how humans, with no knowledge of an environment, explore. The method involves experiments in which participants play an exploration game; in the game, they are asked to explore a multi-floor building and complete a set of tasks within a certain time limit. All the movement and actions of the players are logged and later analyzed for patterns. We then propose a novel way of identifying the role that memory and *non-randomness* plays in human exploration from the data extracted. Thus, the main motivations of this chapter, and therefore its main contributions, are:

1. The development of a novel and scalable game based methodology to study indoor wayfinding behavior.

2. The demonstration of such a scalable method's usefulness in determining how memory influences exploration efficiency and an individual's ability to navigate within an environment.
3. The derivation of an agent-based Markov model that is a good approximation of *typical* human behaviour.

The remainder of this chapter is organized as follows: Section 6.1 first provides an overview of our current understanding of how people gain spatial knowledge and do indoor wayfinding. Following this, the game and experiment are described in Section 6.3. The methodology used for evaluating the results from the game are then explained in Section 6.4. Finally, the results and its implications are discussed in Section 6.5.¹

6.1 Literature Review

Human exploration of indoor environments is a complex process that may depend on many factors including the perception of the environment, the person's existing knowledge and various cultural factors. Before examining existing models of human exploration it is useful to establish a basic knowledge of how spatial knowledge is stored, used and updated in human memory. In the context of this chapter, it is also useful to understand the methodology used in studying human spatial memory as both an inspiration and validation for the methodology adopted here.

6.1.1 Working memory

Since Hebb's work on human memory [150], it has generally been accepted that human memory has a Short Term Memory (STM) component and a Long Term Memory (LTM) component. Baddeley's model of working memory [151], which is a three component model consisting of the central executive, a visual spatial sketchpad and a phonological loop is currently one of the most popular models of the working of human memory. Central to this idea is the concept of *working memory*. Working memory consists of both a visual and a verbal component both of which are limited in capacity.

In order to study the use of memory during wayfinding, Lindberg and Gärling [152] presented people with a difficult task to complete (backward counting), while concurrently performing wayfinding. They discovered that the concurrent task adversely affected wayfinding performance. This supported the notion that navigation may require effective use of limited capacity cognitive sub-systems. However, the relative involvement of different cognitive sub-systems could not be determined

¹This chapter has been submitted to Nature Scientific Reports for review.

through this experiment. Several studies [153, 154] have involved experiments examining the way in which wayfinding memory is stored. These studies concluded that when maps were used for learning, only the visuo-spatial component of memory was used. However, in the real world, they found that all the senses were used together and that verbal, visuo-spatial, temporal, auditory and even olfactory cues were used by the participants. This supports the notion that experiments studying wayfinding without maps should be as immersive as possible to ensure they reflect reality.

Evidence suggests that salient (distinctive) cues are important for place learning. This is especially true as people age and reduce their ability to perceive and process cues [27] or if the way-finders are stressed or under time pressure [26]. This is quite probably due to the fact that both stress and old age can reduce working memory capacity. A decreased working memory capacity implies a smaller amount of environmental information is processed and less information is eventually encoded. This, in turn, implies that only cues that have high perceptual, cognitive or contextual salience are perceived [27] in stressful situations like emergency evacuations that are the main subject of this thesis.

In summary, human working memory consists of different components for processing different types of cues. Experiments show that different cues are used during wayfinding and thus an experiment to study this behavior should be as *immersive* as possible. Also, wayfinding behavior is likely to be different in stressful scenarios due to cognitive constraints in cue processing. Thus, recreation of a stress or time pressure is necessary to recreate egress wayfinding behavior in experiments.

6.1.2 The building blocks of spatial knowledge

There are different *scales* at which locations are stored in the human mind. According to Sholl and Fraone [155], these can broadly be divided into 3 levels: Figural Spaces (Object Sized Space), Room Sized Space (Vista Sized Space), Environment Space (Map Sized Space). The accepted approaches and methodologies employed to study and understand these different scales depend on the scale concerned [155]. In this research, we specifically consider Vista and Map sized spaces.

One of the earliest and most influential works on how humans gain and store knowledge of space was Lynch's *The Image of the City* [132]. He coined the term *mental map* which refers to a person's perception of the world around him. Perhaps the most important contribution of Lynch's paper was the proposal of the fundamental building blocks of a mental map: paths, edges, districts, nodes and landmarks. Paths are routes along which people move, districts are distinct regions, edges define the boundaries between these regions, nodes and landmarks are locations that are points of reference in the mental maps.

How these blocks form a person's image of space was explored by Siegel and White [135] through their experiments on map learning in children. They pro-

posed a hierarchical model with three distinct parts. They found that people first gain knowledge of the *landmarks* in an area, subsequently they learn *routes* connecting these landmarks and finally they gain *survey knowledge*, wherein they have an overall map of the region to the extent that they can determine shortcuts and best routes. They further postulated that adults, despite having more developed cognitive abilities than a child, mirror a similar process in forming their memory of space.

Ishikawa and Montello [156] emphasized that different people have different abilities and techniques for formation of spatial knowledge. Significantly, they found that given repeated exposure to the environment, some people were inherently good and others inherently bad at forming and using spatial knowledge.

Some researchers [131, 156] argue that, contrary to Siegel and White, people's route knowledge and knowledge of space does not significantly improve after first being formed. More interestingly they discovered that certain routes are learned and these routes do not change much; however, inter route connections improve as experience increases.

In summary, while people do have different techniques of forming and storing spatial knowledge, most studies have confirmed that there is a definite pattern in which it is formed, with landmarks and routes playing a key role at the beginning. Furthermore, studies have also shown that this initial knowledge of routes does not change much with time. This implies that the way in which people first explore an environment most likely plays a key role in determining the route they learn and use. This clearly indicates the crucial role that exploration plays in both short-term and long-term way finding behaviour, thus motivating the need for a better understanding of the way in which humans explore environments.

6.1.3 Wayfinding in indoor environments

Literature on human behavior in internal environments, in general, is much more limited than on outdoor wayfinding. Best [157] was the first to identify that the number of choice points, that is, locations where directional changes occurred, was the relevant measure for assessing way finding difficulty, as opposed to simple metric distances.

Weisman [158] defined visual access, degree of architectural differentiation, signs and floor plan configurations as the factors that determine wayfinding difficulty. *Visual Access* which, in essence, refers to the fact that an environment's external structure gives clues to the internal layout and hence visual access to the outside can decrease wayfinding difficulty. The findings of Gärling et al. [159] confirmed the importance of familiarity and visual access. Evans and Pezdek [160] discovered that distinct wall colors reduced wayfinding complexity.

Thorndyke and Hayes-Roth [133] established that map learning can result in survey knowledge developing within 20 minutes. However, if knowledge

is gained through actual daily movement through the environment without the use of maps and compasses then it may take up to 2 years for survey knowledge to be developed. This was further tested by Moeser [131] in their experiments with nursing students in a five-floor hospital building. They discovered that floor plans were never used and that even after 25 months, most nurses don't seem to develop survey knowledge. They rather created an image of the layout of the building within the first month and later experience just extended this in minor ways.

According to Gopal et al. [137], paths with more choice points like intersections or turns are considered more complicated than paths with fewer choice points. This is a natural consequence of the fact that there is a higher chance of error. More interestingly, they state that paths with more turns are perceived to be longer as well.

Kuipers' [134, 136] developed the TOUR model of spatial knowledge processing for *large-scale urban spaces* which was one of the first models of spatial cognition. Part of it was a model for exploration which makes use of existing landmarks. In this model, people explore by maintaining their heading with respect to known landmarks like a tower that can be seen from multiple locations. They return to these locations when required. This exploration methodology was further explored and elaborated in [136]. Here, the authors proposed that there are a small set of paths, which they call *the route skeleton*, which people primarily use for wayfinding. People typically explore along this central route skeleton. This theory was tested through *desktop virtual reality* experiments. However, due to the time required to participate in the experiment and the amount of effort required it was probably both costly and difficult to get many participants. Thus the experiments were conducted only on four participants which, we believe, limits the analysis that can be done on the data.

The importance of structural landmarks like T-junctions in wayfinding was demonstrated by Stankiewicz and Kalia [140] using desktop virtual environments. A slightly modified version of the same desktop virtual environment was used by the authors [139] to study exploration in unknown environments through comparison against an *Ideal-Navigator Model*. They explored the inefficiencies in human navigation in unknown environments by comparing against an ideal navigator, i.e. one who has perfect perceptual processing, perfect map memory and the ideal decision strategy. The agent based analysis discussed in Section 6.5 uses a similar approach. While the experiment structure and the analysis of these experiments are interesting, they were only able to get between three and eight participants for each experiment due to the length of the experiments. Thus there is a risk of sampling bias in their conclusions.

Hölscher et al. [138] conducted experiments to investigate the strategies that people use in exploring multi-floor buildings. The first strategy, discussed earlier, shown by Kuipers et al. [136] stresses on the primacy of a set of central paths or a

route skeleton in way finding. Another strategy, referred to as the *horizontal position strategy*, was used by very few people and generally proved to be inefficient. In this people would try to get to the correct horizontal location first before they try to find the way to the correct floor. The reduced efficiency of this strategy was a consequence of the fact that the experiments were conducted in a building where each floor has a different layout from the next (as is the case in many buildings). The last strategy, which was used by more experienced participants and also proved to be the most efficient, was called the *floor first strategy*. As the name suggests, this strategy involved the person trying to get to the required floor first and then exploring horizontally to find the goal.

O'Neill [161] studied the accuracy of simulated environments in studying wayfinding behavior. Their approach was to examine behavior in a simulated environment and compare it to results from actual experiments. Their findings showed that human behavior in simulated environments was a reasonable approximation of the same real life behavior. Montello et al. [145] had a much more comprehensive analysis of the effect of different sources: maps, virtual environments and real world experience. In general, they observed that the more immersive the environment is, the more likely it is to produce realistic behavior from the participants. However, it is important that immersiveness not be equated with realism. It is unlikely that immersiveness by itself is a necessary or a sufficient criteria to illicit realistic behavior from human beings. A well designed yet abstract setup like [161] is a good counter example. However, a simple point and click setup like the one used in [141] would probably find it difficult to produce realistic behavior from the participants. It is interesting to note that Stankiewicz and Eastman [162] failed to produce this effect of immersiveness producing more realistic behavior. However, this might have been due to the small size of sample data available.

In conclusion, we find that several different kinds of experiments have been conducted ranging from real world, to simulated environments and even virtual reality experiments. Several factors like architectural complexity and visual access have been identified as contributing to wayfinding difficulty. Studies have also revealed that during real world navigation without maps people tend to form an initial cognitive map of the place and this does not improve much over the short term. Thus, the way an unknown environment is first explored is crucial to determining the cognitive map that is developed. There are several theories for strategies that are used during exploration of a large, unknown environment like a multi floor building; of these, a *floor-first* strategy seems to be the most popular and efficient. Some studies have shown that immersiveness does help produce more realistic behavior. However, existing immersive experiments using desktop virtual reality have generally found it difficult to attract more than a few participants. Using serious games for behavior experiments not only overcomes this but has certain other advantages as well that are

discussed in more detail next.

6.2 Benefits of a Game Based Approach

There are four basic advantages that a well designed game based approach can have as an experimental methodology: Immersiveness, Experimental Control, Interestingness and Ubiquity [142].

- **Immersiveness:** Immersiveness or accuracy refers to the fact that game mechanics can be designed to closely reflect reality. As explained earlier, what is important is that the player perceives the artificial environment created as life-like and therefore, makes the same decision in the game as they would in the equivalent real world scenario. For example, a first person environment with realistic but simple movement and turning behavior will likely reproduce more realistic behavior than a point and click game with a bird's eye view.
- **Experimental Control:** The experimental control allowed is a great advantage that a game based approach shares with other traditional desktop virtual reality experiments. Depending on the development environment chosen, the setup of the experiment can be varied at much less cost in terms of time and effort. For example, in most cases, depending on the motivation of the study, it wouldn't take much effort to do A-B testing by varying small elements in the game like the location of a goal provided enough data points are available.
- **Interestingness:** Games are generally much more enjoyable than standard social experiments. This not only encourages repeat participation but it also helps obtain larger participation through recommendations. There are several standard game design fundamentals which have been developed over the past four decades of computer game development that can help create this interestingness.
- **Ubiquity:** As discussed in Section 6.1.3, a limitation of existing desktop virtual reality experiments is the limited data that is available since they generally get fewer than ten participants [136, 139]. While the dataset size was sufficient for some of the conclusions that were made, in several experiments significant conclusions could not be derived. This may be because of the paucity of data. Games can be carefully designed in order to make a repetitive mundane task interesting, which in turn will encourage people to play the game and unknowingly perform the task that is required. Many researchers have collected fascinating statistics as to the number of man-hours spent on gaming around the world. As one example, Jane McGonigal [163] has suggested that the average American 21 year old will have spent 10000 hours of their lives playing games, which is approximately the same number of hours that same person will have spent within

the education system. On May 23rd, 2010 the Google main page was changed to include a version of the classic PacMan game, some estimations are that 4.82 million man hours of productivity were lost in a single day (approximately US\$120 million).

Thus, if designed carefully, a game based experimental methodology can provide several advantages over traditional approaches. In the next section, we present the game that we developed to study how memory influences exploration efficiency and an individual's ability to navigate within an environment.

6.3 Setup of the Experiment

Wayfinding experiments in virtual reality environments consist generally of two parts: Knowledge Acquisition and Task Performance based on this knowledge. Most existing experiments try to control the knowledge acquisition part of wayfinding. A typical example is [154] where the author ensured that all participants received identical stimuli in order to be able to fairly compare their task performance based on this knowledge acquisition. Participants were asked to watch a video rather than actively navigate through the environment. This has been the typical approach to date and this has resulted in there being very little existing literature on how people actually explore environments; we believe that the paths taken during exploration could reveal interesting aspects of human exploration.

In order to test this hypothesis and understand more about the way in which humans explore environments and store spatial knowledge, we created a game that requires both exploration and wayfinding and analyzed how the game was played. While creating the environment, we strived to ensure that it had enough complexity and diversity to be engaging and invite exploration [145, 164].

6.3.1 The game

The premise of the game is that the player has been teleported into an abandoned palace where eleven people have been imprisoned in different locations spread across the three floor environment. The objective of the game is for the player to free the eleven prisoners and subsequently follow instructions to open the main gate to the palace and escape. The palace is a three floor building with 44 rooms. The layout of each floor of the palace is shown in Figure 6.1. A player joining the server is spawned at the location indicated on the map with an X.

During the first few minutes of the game the player is presented with the story line and interactively told how to use the controls and play the game through a set of signboards. They also follow a tutorial which helps them free the first prisoner. Subsequently, the player is tasked to find and free the remaining ten prisoners. The

(a) Ground Floor

(b) First Floor

(c) Second Floor

Figure 6.1: Floor plans of the three floors. The X indicates the starting point. The blue color indicates the prisoner locations. It is a modified version of an existing Minecraft map [165]

locations of the prisons, as shown by the shaded areas in the map, are spread all over the building. This is the first knowledge acquisition phase of the game and we call this the *exploration phase*. This phase requires the player to move around and explore the building. This phase can be reasonably equated to what a new visitor to a building (for example, a shopping mall) experiences.

The next phase, which we call *knowledge testing phase*, starts when the

Figure 6.2: Picture gallery

eleventh prisoner has been freed. During this phase, the cognitive map formed by the player is tested through a series of three tasks, which involve operating switches that were hidden during the exploration phase. By not revealing the nature of the second phase to the player at the beginning of the game and also by hiding the location of the knowledge testing tasks, we ensure that the player does not make a special effort in remembering locations which could artificially alter the cognitive map formed.

When the eleventh prisoner is freed and the exploration phase ends, the player is given instructions to proceed to the gallery room (shown in Fig. 6.2). This room would have been examined by the player during exploration. It is the only room in the palace whose walls are covered with paintings and this makes it reasonably likely that the player will remember this location because of its *perceptual salience* [27].

Once the player locates and presses the switch that is revealed in this room, the player is given instructions to move to the second floor library. There are two important factors for choosing this particular location. Firstly, the library has four entrances and is very likely that the player would have entered this room multiple times during the exploration phase. Secondly, being on the second floor, there are multiple paths to this location from the gallery.

Once this location is found, the player is given the final instruction to proceed to their starting location to find the final switch that will open the main gate to the palace. Again, there are multiple routes to this location some of which are significantly shorter than others. Also, being a starting location and in a somewhat central location, it is likely that it has been frequently visited and should have some *cognitive salience* [27]. We built the experiment inside the popular game Minecraft [166] in order to reduce the development effort and to allow for flexibility in environment

creation.

6.3.2 The Minecraft gaming environment

Minecraft [166] is a Java based multi-platform sandbox construction game. The game involves players creating and destroying various types of blocks in a first-person three-dimensional environment. In the original game, the player controls an avatar that can destroy or create blocks, forming buildings, structures, artwork and even entire cities on multi-player servers or single player worlds across multiple game modes. Players can break any block and build any block, provided he/she has the resources. Figure 6.2 gives a better idea of the look and feel of the game.

In the development of the game described in the previous section, we use an existing plug-in [167] that constrains players so that they can only move around in the environment and interact with doors and switches, that is, elements that are essential for the experiment. A second modification is used [168] to keep a log of the movements and actions of the players and store them in a MySQL server for analysis. The player locations at different times, the time at which each prison was opened and the time taken to complete each task in the testing phase are all recorded for this purpose.

Using Minecraft provided several advantages; being a popular game with standard controls and a first person view, immersiveness and interestingness are much easier to achieve. Another useful consequence of the popularity of minecraft is the large community of developers and the number of Java based plugins that are available for modifying the default Minecraft environment. This made it easier to develop the game and make the necessary modifications to make the experiment as required. However, the biggest advantage of using Minecraft is the ubiquity it offers with over 35 million copies sold. If hosted on a publicly accessible server, it would be theoretically not be difficult to get several hundred, if not thousand participants. Efforts that have been made to do this are presented in Section 8.3. In the present chapter, we present the initial results of the game being played by participants on a local server.

6.3.3 Participants

Fifty young adults (14 women, 36 men, age range: 18-30 years) were recruited with fliers posted on the Nanyang Technological University Campus. Participants were compensated \$ 5 for their participation in the game. It was played by all participants on the same computer in a lab with minimal distractions.

Participants were given five minutes to familiarize themselves with the controls of the first person game, which involved using the keyboard direction keys for movement and the mouse for altering the move and view direction. A mouse click

would allow the player to interact with the environment by either opening doors or freeing prisons. The players were given 45 minutes to complete the game. Of the 50 participants, the data from only 44 participants were used, the remaining six experienced motion sickness from the movement in the first person gaming environment and had to quit playing before the game could be completed. Next, we explain the way in which the data was analyzed.

6.4 Analysis

In this section we present the methods that were used to analyze the data obtained from the game. Figure 6.3 shows the room layout shown in Figure 6.1 as an undirected graph. A directed *movement graph* is calculated for each player, with the edges of the graph indicating the direction and time of movement from one node to the next. The analysis presented in the following sections is performed using these graphs.

6.4.1 Types of exploring agents

As mentioned in the introduction, the motivation of our analysis is to determine how memory influences exploration efficiency and an individual's ability to navigate within an environment. In order to do this, we use an approach similar to the ideal navigator model of Stankiewicz et al [139]. We compare the movement of the players to different kinds of simple exploring agents inhabiting the same environment as the player. By comparing the performance of these simple agents to the actual performance of the player, we determine the ways in which a player is different from or similar to each of these agents.

We compare the *movement graph* of the players to three kinds of exploring agents: Markov agents, unbiased random walkers and agents with perfect m -step memory. Just like the actual players, these agents have no prior knowledge and explore the undirected graph shown in Figure 6.3. Our analysis is done based on four sets of directed movement graphs :

- *Actual Players*: This is the set of 44 player movement graphs i.e. the movement graphs introduced at the beginning of this Section.
- *Unbiased Random Walker*: The next node to which this agent moves is chosen randomly with equal probability. Each iteration of the random walk is performed until the walker covers 100% of the environment. A collection of random walks is obtained until the variance in the radius of gyration of generated movement graphs stabilizes. The radius of gyration is the standard deviation of a walker's position to the center of mass of movement [169]. It gives a measure of the locality of the graph.

Figure 6.3: This figure shows a graphical representation of the room layout in Figure 6.1

- *Agent with perfect m -step Memory*: This is a biased random walker with an m -step memory. It moves exactly like the random walker except that it avoids moving back to any of the m rooms it visited previously. If there is no unvisited room, the agent checks its m -step memory for an unexplored junction. If such a junction exists, it goes back to that point and continues exploring in an unvisited direction. If such a junction does not exist the agent simply chooses, at random, a neighbouring location to the current node with equal probability. Depending on the calculation being performed, the path is generated for a specific number of hops or until a specific coverage is achieved. The coverage is simply calculated as:

$$\text{coverage} = \frac{\text{number of rooms visited}}{\text{total number of rooms}} \times 100 \quad (6.1)$$

30,000 such movement graphs are generated for each value of m for each experiment. This was empirically found out to be a good enough value to ensure minimal variation in the radius of gyration.

- *m^{th} order Markov Agents*: These agents mimic the movement of the average of the ensemble of players assuming they had only an m step memory. The action of an m^{th} order Markov agent at a particular point in the path is a function of the actions of the actual players who had taken the same m steps. This is explained in more detail in Section 6.4.2. The calculations are performed for m from 1 to 13. As with the agent with perfect memory, the movement graphs are generated either for a specific number of hops or until a specific amount of coverage is achieved. Also, as before, 30,000 movement graphs are generated for each value of m .

Simulations of movement were used to generate movement graphs for each of these types of agents and the results of this analysis are presented in Section 6.5. Before this, the next section explains the working of the Markov agents in more detail.

6.4.2 Markov Agents

In this section we present how the movement of the Markov agents are calculated. The chief motivation of this Markovian analysis was to investigate the role that memory plays in the exploration of the environment.

We take an m^{th} order Markov model to represent an m -step memory of the explorer, where steps constitute node (room) visits on the undirected graph (Figure 6.3). One way to speculate on the size of the memory used by a human during exploration is to predict a path of length n from some Markov data of order $m < n$. We hypothesize that *if the movement of the actual players can be predicted using an m^{th} order Markov agent, this implies that humans use a working memory of size m steps during exploration.*

In a general m^{th} order Markov process, the basic idea is that the action at any point of time depends only on the previous m actions. By stating that the process of exploration is an m^{th} order Markov process, we assume that the next node that is visited by a player is only dependent on the previous m steps. This is different from an agent with m -step memory that tries to avoid the previous m nodes. Since the next step is dependent on the actions of players who have visited that same subsequence of m nodes, the Markov model theoretically encapsulates other factors like layout, visibility and so forth and thus unlike the former, has *imperfect recollection*.

Figure 6.4 explains the idea of a Markov agent by contrasting it's behavior with that of an agent with simple m -step memory. In the example layout shown in figure — with the rooms C, D and E connected to corridors A and B — we consider the situation where the exploring agent has moved from A to B to C to B to D and back to B. Assuming a 6-step memory, the simple memory agent would definitely visit room E since all other rooms have been visited in the last 6 steps. However, the Markov agent's action in the same situation depends on the actions of the actual players in this situation. Since two out of the three players visited room E and one out of the three players visited room A, the next step of the Markov agent can either be room E or A with $2/3$ and $1/3$ probability respectively. The following section explains in more detail how the Markov agent movement graphs are calculated.

6.4.2.1 Calculation of Markov data

This section explains how the aforementioned Markov calculations are done. These calculations are performed on the *actual player* dataset D derived from the game (Section 6.4.1). To recap, a directed movement graph is stored for each player. Each

Table 6.1: Summary of symbols and their meaning.

Symbol	Meaning
$Pr(A B)$	Probability of occurrence of event A given event B
X_n	Random variable indicating the location in the n^{th} step
$p_{ij}^{(n)}$	Probability of going from state i to state j in n steps
R	Set of all nodes
r	A particular node
N_r	Set of neighbors of node r
$P^{(n,D)}$	List of all paths of length n taken by individuals in dataset D
$Q^{(n)}$	Random variable representing a path of length n
$q^{(n)}$	A particular path of length n
$\phi(a, L)$	The number of occurrences of element a in list L
$x_i^{(q)}$	The i^{th} location in a particular path q

Figure 6.4: Given the situation in (a), the agent with 6-step memory would move next to room E since it has the list B,D,B,C,B,A as the 6 rooms visited immediately previously. However, the action of the 6th order Markov agent in the same situation is dependent on the actions of the players ((b),(c) & (d)) who happened to be in the same situation. Since one out of the three players who were in the given situation next moved to room A, there is a 1/3 chance that the Markov agent will visit room A next.

node of the directed movement graph corresponds to a node in Figure 6.3 and each edge indicates the movement of the player from one node to the next. Each edge also stores the time point of traversal. The symbols used in this section and their meaning are summarized in Table 6.1.

The m^{th} order Markov probability of visiting a node b from a node a is defined as the probability that the m^{th} node after visiting node a is node b .

Mathematically, it can be stated as:

$$p_{ab}^m = Pr(X_m = b | X_0 = a) \quad (6.2)$$

It is assumed that it is a time homogeneous process, that is,

$$p_{ab}^m = Pr(X_{k+m} = b | X_k = a) \text{ where } k \geq 0 \quad (6.3)$$

To calculate the m^{th} order Markov data using dataset D , we calculate the following:

1. First, for each path of length m , the number of times that path is observed in dataset D is counted.

From this, the following equation can be used to derive p_{ab}^m :

$$p_{ab}^m = \frac{|P_{ab}^{(m,D)}|}{\sum_{x \in R} |P_{ax}^{(m,D)}|} \quad (6.4)$$

2. For each path of length m , the likelihood that a particular path is the result of m steps being taken is simply the ratio of the number of occurrences of the particular path of length m to the total number of paths of length m in the database D ; mathematically:

$$Pr(Q^{(m)} = q^{(m)}) = \frac{\phi(q^{(m)}, P^{(m,D)})}{|P^{(m,D)}|} \quad (6.5)$$

3. Using these, we can calculate the probability of each destination given a particular path of length m by using Bayes Theorem.

For m^{th} -order Markov agent, given a particular path of length m , it is possible to generate the $(m+1)^{th}$ step using the previous m steps, that is, the 1^{st} to m^{th} step. This is done using the destination probabilities calculated in step 3. Following this, the $(m+2)^{th}$ step can be predicted by doing the same calculation using the preceding m steps from 2 to $m+1$. This process can be continued to generate a directed movement graph of n edges using just the m^{th} order Markov data. As mentioned previously, calculations are done using 30,000 paths generated like this. A question that might arise for the discerning reader at this point is whether the dataset is sufficiently large to do such Markov calculations reliably for large values of m . We present further analysis on the validity of this calculation in Appendix B. Next, we look at the results of our analysis.

6.5 Results

From the game, the complete movement traces of 50 players were obtained. This data was analysed by comparing against agents of different types to determine patterns and gain a better understanding of the movements of the players. This section presents the results of four kinds of analysis that was performed. First, we present the results of a simple check of the frequency of visits to each room. Following this, we compare the efficiency of movement of the players against the different agents. Finally, we present some empirical observations from the movement graphs of the actual players.

6.5.1 Room visit frequencies

In our first experiment, we calculate the frequency of visits for each room per player and compare this against the unbiased random walker. This is a simple test to determine if the players have a pattern or strategy in their exploration that is different from a random walk. If there is a significant difference in the number of times a particular floor or area of the building is visited by a player then this will probably be revealed by this comparison against an unbiased random walker.

For each room r we calculate the normalized number of visits α . For any room r , let $f_p(r)$ be the total number of times room r is visited by the the player p and let $f_{rw}(r)$ be the number of times the average random walker visited the same room. We define the normalized number of visits α as:

$$\alpha_x(r) = \frac{f_x(r)}{\sum_{a \in R} f_x(a)} \quad (6.6)$$

This is calculated for both the player and the random walker. Using this the visit ratio, $y_p(r)$, of the actual player is calculated for each room as:

$$y_p(r) = \frac{\alpha_p(r)}{\alpha_{rw}(r)} \quad (6.7)$$

A random walker's frequency of visits to a particular node is purely a function of the topography of an environment i.e. the nodes and their connectivity. Thus, unlike $\alpha_p(r)$, $y_p(r)$ would not contain the effects of the degree of a node. This means that if a room r_1 has lower value of lower $y_p(r)$ than r_2 , this is not because of the room having lesser connectivity. The average $y(r)$ over the ensemble of player data was then calculated. This is illustrated in Figure 6.5. This figure shows the value of $y(r)$ of each room as a scaled version of Figure 6.3. The red color indicates those rooms that have 5% more visits than a random walker and the green color indicates those that have 5% less visits than a random walker; the node is colored white otherwise. The diameter of each node in this graph is scaled to $y_r \times \text{unscaled diameter}$.

Therefore, a white color indicates that the normalized number of visits

Figure 6.5: Figure 6.3 scaled by normalized number of visits as described in Equation 6.7. Red color indicates a y value of greater than 1.05 and the green color indicates a value of less than 0.95. The diameter of each node in this graph is scaled to $y_r \times (\text{unscaled diameter})$.

in both the *random walker* dataset and in the *actual player* dataset are within 5% of each other. A value greater than 1 indicates that players visited the room more than the random walker and a smaller value indicates the opposite. It is hard to discern any pattern in Figure 6.5, other than the fact that there is a marked difference in the visits by the players and visits by the random walker. However, it can be seen that the average number of visits on the second floor was higher than the number of visits on the first floor which was more than the number of visits on the ground floor.

Despite a random walker not differentiating between staircases and other links, a random walker might have more visits to one floor than another purely because there are fewer inter floor connections than other links. For example, floor 3 has only one staircase that leads the random walker out of the third floor. To confirm that the inter floor differences for the actual player is not because of this, we normalize Equation 6.6 to the number of visits on the floor rather than the total number of visits. On doing this, if the graph is different from Figure 6.5, it would seem that unlike a random walker, a player differentiates between a simple link between rooms or corridors and a staircase which is a link between floors. Figure 6.6 which is the floor normalized version of Figure 6.5 does clearly show this.

Figure 6.5 and 6.6 provide a possible validation of a variation of the floor first strategy [138] used for exploration. The strategy in the original paper was for wayfinding with a particular goal, but here it seems to be being used for exploration (which is wayfinding without a definite goal). The players seem to consider each floor as a separate entity and are generally reluctant to take the staircase. This might also be because the process of separating each floor helps in bringing some organization and structure to the confusing room layout and the process of exploration (that is, completely explore one level before the next).

There is a chance that this aversion to taking staircases is a consequence of the structure of the game environment or the controls in the game. This would happen if the staircases were difficult to find or climbing staircases required different controls from moving through doors or walking through corridors. However, this is not the case in the designed game. The player just has to press the forward button to move in the environment regardless of whether it is a staircase or a corridor or a room. Also, the staircases are clearly visible from the neighboring nodes with sign boards further confirming their location.

The hypothesis of the existence of this floor first strategy is further strengthened by the low visit frequencies to Block B on the second floor 6.6. There are four staircases that take a player from floor 1 to floor 2: two of these lead to Block A, one to Block C and one to Block B. Block A and C are directly connected by a corridor whereas the only paths to Block B from the same floor are through rooms DB2 and The Lounge in Block A and C respectively. This means that, unlike Block A and C, Block B is not accessible via any direct corridor from the same floor. Since people

Figure 6.6: Figure 6.5 normalized to floor instead of the total number of visits. The fact that this graph is different from Figure 6.5 indicates that, unlike a random walker, a player differentiates between a simple link between rooms or corridors and a staircase which links two floors.

Figure 6.7: This figure shows a standard error plot of the average coverage after 253 hops as a function of memory size. The low values of standard error on the agent paths are because these calculations are calculated over several thousand paths that are required for the radius of gyration to stabilize. A value of 253 hops was taken because this was the average number of hops taken by a player.

show a clear inclination to exploring through corridors as shown in Section 6.5.4 of this chapter, the only obvious way to access Block B is by going down a floor. The fact that Block B has fewer visits seems to suggest that people resist going down a floor. Thus further strengthening the hypothesis of a floor first strategy in exploration.

6.5.2 Expected coverage given number of hops

The average coverage after a given number of hops gives an estimate of the efficiency and effectiveness of exploration. On average each player took 252.6 ± 7.5 hops during their exploration phase. The coverage achieved by the other types of agents after 253 hops was calculated by generating the required set of movement graphs as explained in Section 6.4.1. Figure 6.7 shows the results of this calculation.

The figure seems to indicate that even a second order Markov agent, that is, one whose next position is only dependent on its current and previous position, performs much better than an unbiased random walker. It also seems to indicate that after 253 hops the performance of the actual players is much better than both the Markov agent and an agent with an m -step memory. This is not surprising since it is likely that when nearing 253 hops, the long term memory of the player also has a major

influence. As discussed in the literature review, in the longer term, the player would probably have formed a route or some sort of survey knowledge; this may include the structure of the building, routes and short cuts and, in general, means there is likely more structure to the mental map. The fact that the Markov agent performs worse than the agent with memory regardless of the value of m agrees somewhat with this conclusion. However, this could also be because the Markov agent has the same errors as the collective human memory - whereas the m -step agent has perfect memory (Section 6.4.2).

6.5.3 Expected hops given coverage

We also calculate the minimum number of hops required to obtain a given coverage. Unlike coverage, a hop count captures the dynamics of room revisits. The average final coverage for a player after the exploration phase of the game is $(89 \pm 1)\%$ as shown in Figure 6.7. We first calculate the minimum number of hops required by the different agents to obtain this coverage; this is shown in Figure 6.8a. The graph shows the same pattern as in Figure 6.7. The only difference is that the number of hops required by the Markov agents increases again for large memory steps. This is probably because rather than going to new nodes the Markov agent revisits old nodes; this results in the coverage not increasing.

As mentioned previously the reason for the patterns observed might be effect of long term memory. In order to test this, the same the expected number of hops for 50% coverage was also determined and Figure 6.8b was obtained. The magnitude of the difference between hops required for 50% and 88% shows a non linear increase, indicating that exploration becomes progressively more difficult. The figure also shows that agents with a memory of 5 or more steps seem to perform at the same level or better than humans. It is interesting to note that the performance of Markov agents is worse than agents with a simple m -step memory. Again, this is probably due to the imperfect nature of the short term human memory on which it is based (Section 6.4.2). The gap in performance between the Markov agent dataset and the actual player dataset is quite narrow at $m = 7$ to 9. This indicates that the room visited at any point can be reasonably predicted from the previous 6-8 rooms during this early phase of exploration.

6.5.4 Empirical analysis

In this section we conduct an empirical and qualitative analysis of the *actual player dataset*, i.e., the actions of the players at different locations. This analysis is intended to reveal the existence of patterns in exploration like definite decision points and the importance of cues in recognition and memory which were not discernible from the movement graphs discussed previously.

(a) Hops for 88% coverage.

(b) Hops for 50% coverage

Figure 6.8: This figure is a standard error plot of the minimum number of hops required for obtaining given coverage and shows this as a function of memory size. As with Fig. 6.7, the low values of standard error on the agent paths are because these calculations are calculated over several thousand paths that are required for the radius of gyration to stabilize.

Figure 6.9: This chart summarizes the behavior at corridors of different types. Less than 10% of the players make a conscious decision (i.e. change their direction) on a straight corridor, a staircase, or even a simple corner. More interestingly, a little less than 25% go back instead of opening a clearly visible door in front of them. However, when this next door is not clearly visible when they enter, there is 40% of the player just going back through their entrance. This indicates that there are definite decision points during exploration.

6.5.4.1 Existence of decision points

Figure 6.9 illustrates the decisions of people at different types of rooms and corridors, where it is possible for them to make a decision. At certain locations, such as *simple corridors* that have no rooms on the side (that is, they are simply connections between two areas), staircases and simple corners, the only decision that a player can make is whether to move forward or turn back. Turning back would require a conscious decision by the player. A pure random walker would have an equal chance of going back or forward. As shown in Figure 6.9, the data reveals that players generally don't change their mind.

What is more interesting is the behavior of people at rooms which have just two doors. The data seems to reveal that if the opposite door is clearly visible from one door, then the room is used by the player almost exactly like a *simple corridor*, though there is a slightly higher chance of turning back. However, if the opposite door is not clearly visible when the person enters i.e. it is at some angle to the view when the player enters or there is some furniture partially blocking the view to that door, then there is more than 40% chance of the person simply going back.

There is an argument to be made that this reluctance to move towards

Figure 6.10: The layout of the area under consideration for analysing location recognition and memory.

a slightly less visible door is a result of the input controls (keyboard and mouse) of the game. However, we believe this is not the case because the data that was used was observed from players who did not find it difficult to turn. Moreover, only the data of players who managed to complete the game within the time limit was included. As doing this required exploring a complicated 3 floor environment with 44 rooms and plenty of turns, the players had to be able to use the movement and turning controls in a reliable and natural manner.

In order to demonstrate the ability of the Markov model to encapsulate factors other than memory that influence movement, we also plot the same behaviour for a 7th-order Markov agent, a simple agent with 7-step memory and an unbiased random walker. As expected, the Markov agents reflect the same patterns as the actual player while the behaviour of the other agents are completely different: the random walker makes no difference between going forward or back and the agent with memory never goes back.

6.5.4.2 Location recognition and memory

In the game environment, there exists a corridor that seems to reveal an interesting aspect of memory and exploration. The layout of this corridor is shown in Figure 6.10. The corridor labeled *Dorm Corridor* is interesting because it is connected to the main Block A corridor only at one end². The two rooms on this corridor (D1 and D2) do not have a checkpoint, a staircase, or any connections that make it at all relevant to the player. However, it lies on a commonly used corridor (marked Block A Main Corridor) and is used by most players at least once. In an ideal scenario, players

²With reference to the previous discussion on decision points, it is important to remember that this is not a *simple corridor* because it has doors on the side and is thus a decision point.

Figure 6.11: This figure shows the movement of the players at the junction of Block A Main Corridor and the Dorm Corridor during the Knowledge Testing Phase. As can be clearly seen, players keep revisiting the corridor despite having visited it before

would remember this fact and never visit *Dorm Corridor* after the first visit to the junction of *Block A Main Corridor* and *Dorm Corridor*. The actual movement of the players at this junction is shown in Figure 6.11. Surprisingly, the figure shows that during the course of the game, regardless of the number of times the junction is visited, players turn into the *Dorm Corridor* more often than not.

At first this leads to the conclusion that the players never learn and have no memory. However, a similar analysis of movement after entering *Dorm Corridor* indicates that this isn't the case. As shown in Figure 6.12 about 80% of the players head right back to the junction after entering this corridor. In order to understand why this happens we can compare the player behavior against the Markov agent. We see that at the junction, the actual player movement is similar to the Markov agent movement suggesting that the movement is primarily governed by short term memory. However, in the Dorm Corridor the actual players behavior is completely different from the Markov Agent after the first visit, suggesting that long term memory is being triggered by some factors at this point. This probably indicates that the context given by the location of signs and doors in the corridor helps the player remember the corridor, its location and its use.

This could theoretically be verified through the continuous measurements of mouse movement trajectories of players at these points as in [170, 171]. However, this constant sampling and recording may have consequences on the frame-rate and playability of the game. Also, analyzing such information is complicated by the difficulty of matching mouse location with what exactly is seen on the screen as the player is constantly moving around a virtual world. However, this is definitely an avenue for future analysis.

Figure 6.12: This figure shows the behavior of the players after entering *Dorm Corridor* during the Knowledge Testing Phase. At this point, most players remember the corridor and head back to the Main Corridor without exploring

6.6 Summary

The spatial knowledge of an evacuee determines the path that is taken by him/her during evacuation. Traditionally, agent based egress models either assume complete spatial knowledge or have simple eidetic memories where, once perceived, locations are never forgotten. Thus, a need for more accurate modeling of spatial cognition and exploration of unknown indoor environments was identified. In order to enable this, in this chapter, a novel game-based methodology was developed and used for the experimental investigation of indoor navigation and exploration.

The novel Markovian analysis of the player's movement in the game revealed a number of significant findings. Firstly, we showed that a simple memory model, with a depth of between 6-8, is sufficient to approximate a 'human level' of exploration efficiency. This was consistent in two measures of exploration efficiency: total coverage from a fixed number of hops and the number of hops required to obtain a fixed coverage. The memory depth of 6-8 seems to be consistent with well known studies of human memory capacity. The experiments also highlighted the importance of junctions in the exploration process, in particular, decisions (that is, changing course) seem to almost exclusively occur at junctions. Explorers also try to reduce the number of decisions they have to make by proceeding to the next clearly visible room or corridor (assuming only one such room is visible). Furthermore, the results showed that people seem to explore environments using a floor first strategy, where they are reluctant to move to a different floor until they have finished exploring the current one. Finally, it was shown that easily recognizable locations probably help individuals improve exploration efficiency by enabling individuals to effectively remove sub-graphs of the room network from their cognitive map.

Several of these observations could be made only through the analysis of the movement graphs and the Markovian analysis on the data. It would have been difficult to obtain this amount of data through the traditional experimental methodologies without significant cost in terms of time and effort. Thus, a game based analysis is not only useful in making significant observations about human behavior but in also helps open up new possibilities of research.

Besides revealing several insights into how unknown indoor environments are explored, the research in this chapter also revealed new models that could form an excellent starting point for improving the exploration models used in existing crowd simulations. The next step in a behavior model of egress is the simulation of the actual movement of the evacuee towards a chosen destination based on the evacuation plan formed. There exist several different models for modeling this movement; however, a consensus on the most realistic movement model or even a methodology for comparing models does not exist yet. This problem is explored further in the next chapter.

Chapter 7

Quantitative Comparison Between Motion Planning Systems for Evacuation Planning and Evaluation

The previous chapters have introduced methods that enable the modelling of more realistic perception, pre-evacuation and subjective knowledge based evacuation planning. Motion planning, which deals with how exactly an agent moves from one location to another, is a critical and at the core of not just the simulations used in the previous chapters, but also of all crowd simulations, egress or otherwise. As a result, there have been several models for motion planning that have been developed over the years. However, there is little consensus on which is the best model to use and, more importantly, no method for determining this. This chapter introduces a methodology for comparing different models of crowds and identifies a key metric that can be used in differentiating the different models ¹.

As discussed in Chapter 2, the earliest attempts to simulate crowds generally adopted a macroscopic approach [67, 74], where there is no explicit notion of an individual. Later, with increasing computational resources and with availability of observational data on an individual level [129, 172, 173], modelers were able to develop microscopic approaches [7, 10] for application in areas where it was necessary to model and analyze individuals in the crowd. For example, in a simulation of evacuation, knowledge of the movement of crowds could reveal methods to improve crowd flow and evacuation speed.

The advocacy of simulation-based analysis has become increasingly common over the last decade. Some well known applications include analysis of the yearly Muslim Hajj [174], or more recently the Love Parade disaster, Germany 2010 [2]. In the case of the latter, models and expertise had been used to guarantee the safety of the event, only for unforeseen circumstances to result in the deaths of 21 individuals.

¹This chapter was published in the European Physical Journal B [108] and featured on Science Daily (<http://www.sciencedaily.com/releases/2014/02/140207114150.htm>)

Clearly these real world examples emphasize the critical role that crowd modelling plays in safety preparation and planning, this in turn emphasizes the need for understanding the model dynamics, limitations and similarities. The extent to which these models are capable of accurately predicting the motion of a crowd is therefore critical for planning and safety.

Ideally these models should be validated against real-world data from all scenarios, for all cultures and for all varieties of crowd composition. Unfortunately real world data regarding egress or emergency situations is limited, often incomplete (the initial conditions are hard to know) and certainly not controlled. Even in the rare circumstances where data is available this often describes measurements and phenomena at the macro-scale, e.g., flow, density, average speed, etc. Because these models of individuals exhibit emergent behaviors at the mesoscopic and macroscopic scales, it is in general hard to tell whether the dynamics at the individual level are correct. The best researchers can realistically hope for is to collect microscopic data for a single scenario to calibrate the model. An alternative approach to the data-intensive one is to instead look quantitatively at the fundamental dynamics of the models and understand where these models differ, at the same time identifying common aspects of the models. This will aid in understanding fundamental aspects of all crowd models and offer insight into real-world dynamics of human crowds.

While much of the existing research does offer basic comparison of proposed models with existing models to demonstrate their usefulness, there is no existing quantitative comparison of the differences and similarities of these models to tell which is most accurate. The objective of this chapter is to demonstrate a methodology for quantitative comparison of simulation models in a simple egress simulation. For this, three popular models that are structurally very different are chosen. Since our focus is on the comparison methodology, specific versions of the models are not important, in particular recent modifications and enhancements that are supposed to produce more realistic crowd behaviors.

The contribution of this chapter is then three-fold: firstly, the analysis and comparison of these models provides interesting insight into the consequence of adopting each in particular forms of crowd or egress simulation. Secondly, the systematic approach we describe could be used in the future to compare further models and develop a standard method of comparison for crowd simulation. Finally, the chapter conclusion identifies a single measurable metric that is most effective in distinguishing the behaviour of the models.

The remainder of this chapter is organised as follows. The models that are considered in the comparison are first described in Section 7.1; Section 7.2 describes the experiments that were conducted and the methodology for analysis. Significant observations from the simulations are presented and analyzed in Section 7.3. Finally, Section 7.4 concludes the chapter.

7.1 Motion Planning Models

One of the earliest works in individual-based motion planning was Craig Reynolds's model for coordinated animal motion such as bird flocks and fish schools [7]. Since then there have been many different approaches ranging from approaches using magnetic [119] or coulombic forces [80] to synthetic vision based models [120]. Section 4.2 provided a brief overview of some of the different approaches used for motion planning. The different approaches that are used for motion planning can broadly be categorized into three kinds: Cellular Automata (CA) based models, force based models and velocity based models. Unlike force based and velocity based models, Cellular Automata based models, introduced in Section 2.3.2.3, consider space to be discrete. Lattice Gas Models are typical CA models that are popular for simple scenarios like analyzing evacuation dynamics from a single room. Force based models [10, 80, 119] make use of fictional attractive and repulsive forces to model how agents try to move towards a goal while avoiding collisions with static and dynamic obstacles. Helbing's Social Force Model [10] is, by far, the most popular of these models because of its intuitiveness and the ease with which more complex behaviors like panic and grouping behavior can be simulated. More recently, there have been several velocity-based approaches to motion planning. Of these, the Reciprocal Velocity Obstacle Model [103, 121, 122] seems to be gaining in popularity because of its computational efficiency and its ability to produce smoother results in certain standard scenarios. Thus, the three models chosen for comparison in this study are the Lattice Gas Model, the Social Force Model and the Reciprocal Velocity Obstacle Model.

In this section, the three individual (or microscopic) models that are compared in this chapter are described. We use the general term agent to refer to individuals within each of the models. The lattice gas model is a probabilistic approach where the future location of an agent is probabilistically determined based on the current configuration of its neighborhood. This implies that the same initial configuration can produce different results during different runs. On the other hand, in the social force model and reciprocal velocity obstacle model, agents find their ways to their destinations via deterministic collision-avoiding calculations. As a result these models produce the same result on multiple runs for a particular initial configuration. The following sections give an overview of each of how each of these models work.

7.1.1 Lattice gas model

A Cellular Automata (CA) model is one in which space and time are discrete. Furthermore, the state space of a CA model is also discrete and finite. In each time step the values of all cells are updated synchronously based on the values of cells in their neighborhoods. Depending on the type of neighborhood (i.e., von Neumann, Moore), and the type of lattice (triangular, square, hexagonal, etc.), the exact number

Figure 7.1: Four out of eight possible configurations of a walker on the square lattice moving towards an exit in the direction $\vec{e}$ shown: (a) unobstructed walker, (b) walker obstructed to the left, (c) walker obstructed to the left and top, and (d) completely obstructed walker.

of cells in the neighborhood of a given cell can vary [77]. Lattice gas models are CA models that make use of a discretized version of the Boltzmann transport equation to model motion [66, 81, 82]. Advances in this modelling area include the Extended Floor Field Model [83] in which agents interact through virtual traces that act like the pheromones in chemotaxis, and the SWARM information model [84] which uses multiple floor fields to model transmission of knowledge between agents. In this study the model by Tajima and Nagatani [175] is used. In contrast to later CA models with additional features like bi-directional movement [176] and vision impairment during evacuation [66], this original model is simple, and adequate for our model comparison purpose.

A square grid is used, and each cell can be occupied by at most one person. The person performs a random walk biased towards the single exit in the room. For example, in Fig. 7.1(a), the probability that the person takes a step in the $+y$ direction is $P_y = (1-D)/3 + De_y/(e_x+e_y)$, while the probability that it takes a step in the $-x$ direction is $P_{-x} = (1-D)/3 + De_x/(e_x+e_y)$, since the intended direction $\vec{e}$ has positive projections along the $+y$ and $-x$ directions. Along the $+x$ direction, for which $\vec{e}$'s projection is negative, the probability that the person takes a step in the $+x$ direction is $P_x = (1-D)/3$. Here we see that the random walk probabilities contain an unbiased component $(1-D)/n$, which is the same for all n permissible directions (in the Tajima-Nagatani model, the $-y$ direction is not permitted), as well as a biased component $D = 0.7$ that favours movement towards the exit.

7.1.2 Social force model

The social force model is one of the most popular models for motion planning in crowds [123, 177, 178]. This model is based on the idea that pedestrians move in response to fictitious attractive or repulsive *social forces* produced by obstacles and other pedestrians. Over the years, several extensions have been made to the social forces model like the modelling of grouping behavior [123]. In this chapter, the Helbing-Molnár-Farkas-Vicsek (HMFV) social force model [10] is used, which is adapted from the original model introduced by Helbing and Molnar in 1995 [10]. In the original model

$$m_i \frac{d\vec{v}_i}{dt} = \vec{f}_i + \sum_{j(\neq i)} \vec{f}_{ij} + \sum_W \vec{f}_{iW} \quad (7.1)$$

three types of forces act on a given agent i with mass m_i , instantaneous position $\vec{r}_i(t)$, and instantaneous velocity $\vec{v}_i(t)$. The first is a restoring force

$$\vec{f}_i = -m_i \frac{\vec{v}_i - \vec{v}_0}{\tau_i} \quad (7.2)$$

that steers the agent towards the desired velocity $\vec{v}_0$ at a rate determined by the characteristic time $\tau_i = 1$. Here

$$\vec{v}_0 = (1 - p)V_0\vec{e}_i(t) + p\langle\vec{v}_j\rangle_i, \quad (7.3)$$

where V_0 is the preferred speed, $\vec{e}_i$ is a vector that points towards the exit, and $(1 - p)$ is the weight given to this desired velocity. With a weight of p , agent i also adapts to the average velocity $\langle\vec{v}_j\rangle_i$ in its neighborhood. When p is small, agent i moves more along its intended direction $\vec{e}_i(t)$, whereas if p is large, agent i tends to follow where its neighbors are going. We can therefore tune p from $p \approx 0$ (self-directed normal egress) to $p \approx 1$ (panic-driven herding during emergency evacuations). In this chapter, $p = 0.2$ was used to simulate an emergency evacuation situation.

The second is a repulsive force

$$\begin{aligned} \vec{f}_{ij} = \{ & Ae^{(R_{ij}-d_{ij})/B} + k\eta(R_{ij} - d_{ij})\} \vec{n}_{ij} \\ & + \kappa\eta(R_{ij} - d_{ij})\Delta v_{ji}^t \vec{t}_{ij}, \end{aligned} \quad (7.4)$$

where

$$\eta(x) = \begin{cases} x, & \text{if } x \geq 0; \\ 0, & \text{if } x < 0 \end{cases} \quad (7.5)$$

that mimics the *psychological tendency* of agents i and j to move away from each other if they are too close. Here $R_{ij} = R_i + R_j$ is the sum of radii of the two agents, $d_{ij} = |\vec{r}_i - \vec{r}_j|$ is the physical distance between the two agents, and $\vec{n}_{ij} = (n_{ij}^1, n_{ij}^2) = (\vec{r}_i - \vec{r}_j)/d_{ij}$ is the unit vector pointing from agent j to agent i . Additionally, $A = 2000$ N and $B = 0.08$ m are the repulsion coefficient and the fall-off length of interacting

agents respectively [179]. Helbing and Molnar also found it necessary to introduce two other terms in the interaction force when agents i and j are in contact with each other, i.e. $d_{ij} < R_{ij}$. This counteracting body compression term $k\eta(R_{ij} - d_{ij})\vec{n}_{ij}$ and sliding friction term $\kappa\eta(r_{ij} - d_{ij})\Delta v_{ji}^t \vec{t}_{ij}$ are crucial for getting realistic behaviors of panicking crowds. Here $\vec{t}_{ij} = (-n_{ij}^2, n_{ij}^1)$ is the tangential direction and $\Delta v_{ji}^t = (\vec{v}_j - \vec{v}_i) \cdot \vec{t}_{ij}$ is the tangential velocity difference.

Finally, in the third force

$$\begin{aligned} \vec{f}_{iW} = & \{A_i e^{(R_i - d_{iW})/B_i} + k\eta(R_i - d_{iW})\} \vec{n}_{iW} \\ & - \kappa\eta(r_i - d_{iW})(\vec{v}_i \cdot \vec{t}_{iW})\vec{t}_{iW}, \end{aligned} \quad (7.6)$$

the first and second terms repels agent i from a wall that it is d_{iW} away from, while the third term (which is negative) is introduced to mimic the observation that people move faster near walls when they are in crowded situations. Here, $\vec{n}_{iW}$ is the normal vector of the wall, and $\vec{t}_{iW}$ is the tangent vector of the wall. Also, $k = 12,000 \text{ kg/s}^2$ and $\kappa = 24,000 \text{ kg/ms}$ are respectively the body force constant and the sliding friction force constant used.

7.1.3 Reciprocal velocity obstacles model

The third model we consider is the reciprocal velocity obstacles (RVO) model. In this model, the time to the next collision is calculated based on the relative velocities of N agents. Each agent then changes its velocity $\{\vec{v}_i\}_{i=1}^N$ to maximize this time to collision. By continuously updating the velocities in this manner, collisions are avoided. This algorithm was first proposed by Fiorini et al. [180], but was first used in multi agent systems in [124]. Since then there have been several modifications and improvements to the algorithm [12, 102, 103, 181], although the underlying idea remained the same. In this chapter, the RVO2 model introduced by Guy et al. [103] is used, where the collision avoidance computation is based on computational geometry and linear programming.

Given a preferred velocity, RVO2 helps an agent find the velocity closest to the preferred velocity that will enable it to avoid collisions with all other agents. This is done by determining for each neighboring dynamic and static obstacle an *Optimal Reciprocal Collision Avoidance line (ORCA line)*. Each ORCA line determines the half plane in the velocity plane where the agent's velocity should lie to ensure that no collision occurs with that particular obstacle for the next τ seconds. τ is a parameter that specifies the number of seconds for which the chosen velocity should avoid a collision. From the set of half planes thus obtained, the optimal collision-avoiding velocity can be determined as the velocity within the permissible region for all the half planes that is closest to the preferred velocity. This can be determined efficiently using linear programming.

Besides the speed and efficiency of the algorithm, another strong appeal of this model is that we need only set one parameter τ . For the experiments in this chapter two values of τ user are: 0.5 seconds or 10 simulation time steps for avoiding dynamic agents and 0.05 seconds or 1 simulation time step for avoiding static obstacles. The value of τ translates in practical terms to how early a walker tries to avoid a collision with an obstacle.

7.2 Methodology

In this section, the simulations that were carried out and the analysis that was done is explained. Java 6 and the MASON simulation framework [182] were used for creating and running the simulations. The simulated scenario consists of agents evacuating from a simple rectangular room with a single exit. The dimensions of the room are shown in Fig. 7.2. Evacuation of the room was simulated with $N = 50, 100, 150, 200, 300, 500$ and 1000 agents whose locations were randomly distributed within the room. For each N , 100 simulations were run for each model, and the positions and velocities of all agents at all time steps were collected. For meaningful comparison between the three models, it is assumed that each agent is a perfect circle with radius $R = 15$ cm [40], that has a preferred speed of 1.3 ms^{-1} , and a maximum speed of 2.6 ms^{-1} [183]. For the social force model all agents are assumed to have the same mass of 60 kg.

Where possible, the same settings were used for all three models to ensure a fair comparison. The values used for the remaining parameters required for the models were given in Section 7.1 along with the respective models. However, to the best of our knowledge, only the social force model has been calibrated against real world data [184]. The other two models are not calibrated beyond choosing a reasonable and believable model and the values chosen are taken such that they produce empirically believable results.

7.2.1 Measurements

Depending on the context and motivation of the model there are several metrics that are generally used. For CA models that are generally used for studying macroscopic patterns, it is common to measure macro variables like the mean evacuation time [66] and flow rate [175]. For models like social force where the motivation is to study interactions at a more granular level, it is common to make empirical and qualitative observations like the lane effect [10]. Sometimes quantitative metrics like density-dependence of the flow or velocity are also measured [185]. From a graphics perspective, where performance is key, computation time [120], frame rate and run time per frame are generally measured against the number of walkers [122].

Figure 7.2: The environment setup used for the simulations in this chapter.

In this chapter, two main classes of measurements are performed: time-based and distance-based measurements. For time-based measurements, the evacuation time distributions for different number of agents was first measured. To better understand the stages in the evacuation, the room is divided into six different zones (see Fig. 7.3) to measure the evacuation time sub-distribution. The outer radii of Zones 1, 2, 3, 4, 5 and 6 are 5 m, 10 m, 15 m, 20 m, 25 m and 30 m respectively. The flow rate at the exit as a function of time is also measured for each model.

For distance-based measurements, the trajectories of all agents were traced to obtain the distribution of total distance travelled going from the initial position to the exit. Here, the total distance travelled D_i by agent i is just the sum of its displacements $\sum_{t=0}^{T_i} \|\vec{r}_i(t+1) - \vec{r}_i(t)\|$, where T_i is the evacuated time of agent i . The room is divided into 100×100 cells, and the number of agents passing through each cell is counted as the simulation progresses. This is then visualized as a heat map. Finally, the ratio of the total distance travelled D_i by agent i to the minimum distance $D_{\min}$ it would cover if it evacuated along a straight line is measured (For lattice model, $D_{\min}$ is the Manhattan distance between initial position of agent and exit). This in some sense quantifies the ‘inconvenience’, I_i experienced by the agent i during the evacuation.

$$I_i = \frac{D_i}{D_{\min}} \quad (7.7)$$

Figure 7.3: In the simulations, agents are grouped into six zones based on their initial distances from the exit. The outer radii of Zones 1, 2, 3, 4, 5 and 6 are 5 m, 10 m, 15 m, 20 m, 25 m and 30 m respectively.

7.2.2 DISTATIS

After making six different measurements on three models, it was discovered that when we look at the flow rates, the lattice gas model is more similar to the RVO2 model. However, if we look at the total distance travelled, the social force model is more similar to the RVO2 model. If we look at the zoned evacuation times, we find that the three models have very different distributions (see Section 7.3.1). Since we do not know *a priori* which measurements best discriminate the three models (or for that matter, best discriminate models from the real world), we want to be able to compare the three models quantitatively, incorporating information from all six measurements.

To do so, we compute the *Jensen-Shannon divergence* [186]

$$D_{JS}[f_{k\mu}, f_{k\nu}] = H[f_k] - \frac{1}{2}H[f_{k\mu}] - \frac{1}{2}H[f_{k\nu}] \geq 0 \quad (7.8)$$

between the distributions $f_{k\mu}(z)$ and $f_{k\nu}(z)$ for measurement k of models μ and ν . Here, z is a continuous variable like evacuation time, distance travelled, or inconvenience,

$$H[f] = - \int dz f(z) \ln f(z) \quad (7.9)$$

is the *Shannon information function*, and

$$f_k(z) = \frac{1}{2} [f_{k\mu}(z) + f_{k\nu}(z)]. \quad (7.10)$$

If $f_{k\mu}(z)$ and $f_{k\nu}(z)$ are highly similar, we will get $D_{JS}[f_{k\mu}, f_{k\nu}] \approx 0$, whereas if $f_{k\mu}(z)$ and $f_{k\nu}(z)$ are very different, $D_{JS}[f_{k\mu}, f_{k\nu}] \gg 0$, i.e. the Jensen-Shannon divergence qualifies as a distance metric. In this way, we obtain six 3×3 distance matrices $\vec{D}_k$.

Then, we use the *DISTATIS* method [187] for analyzing multiple distance matrices. This is a generalization of the method of principal component analysis (PCA). STATIS is an acronym for the French expression ‘Structuration des Tableaux à Trois Indices de la Statistique’, which roughly means ‘structuring three way statistical tables’. The difference between a straightforward application of the PCA method and the DISTATIS method is shown in Fig. 7.4. Instead of six different PCAs for the distance matrices obtained from the six different variables measured, and inevitably ending up with conflicting conclusions on which models are more similar, we standardize the six distance matrices, and ask which variables are more similar to each other.

The idea behind DISTATIS is that the data points are similarly clustered in two variables, if their standardized distance matrices in these two variables are similar to each other. Therefore, if we compute the cross correlations between variables, we will discover which variables give more similar outcomes to which other variables through a PCA. Components of the first eigenvector obtained from this PCA tells us how important each variable is. By weighting the distance matrices of each variable by its component in the first eigenvector, we construct a *compromise matrix* whose matrix elements give us the most reliable similarity/dissimilarity between data points. A final PCA of this compromise matrix then gives the most reliable groups of data points. In Fig. 7.12, the six independent PCA results is shown superimposed on the PCA of the compromise matrix.

7.3 Results and Discussion

The six measurements on the simulations are the (1) evacuation time, (2) zoned evacuation time, (3) passage density map, (4) total distance traveled, (5) inconvenience and (6) flow rate.

Here the results for (1) evacuation time (Fig. 7.5) and (2) zoned evacuation time (Fig. 7.6), (3) passage density map (Fig. 7.7), (4) total distance traveled (Fig. 7.8), (5) inconvenience (Fig. 7.9) and finally flow rate (Fig. 7.10) for all three models, with different number of agents are shown.

Figure 7.4: In many classification problems, we can define different distance metrics for the same set of data points to measure how dissimilar they are to each other. In this figure, three different distance matrices $D1$, $D2$, and $D3$ are shown. If principal component analysis (PCA) is performed on them separately (PCA1 , PCA2 , PCA3), we are not likely to agree on which data points are more similar to each other. In the DISTATIS method, we first standardize the distance matrices by centering and rescaling, to get $S1$, $S2$, and $S3$. If two data points are indeed more similar to each other than they are with the third, then we expect this similarity structure to be reflected in $S1$, $S2$, and $S3$, i.e. two standardized distance matrices would have very similar patterns of matrix elements. With this in mind, we reshape the matrices $S1$, $S2$, and $S3$ to make them into vectors of distances. We then compute the cross correlations between these vectors, before performing PCA on the correlation matrix C . This PCA tells us which distance matrices agree with each other more, and which less. The result is a consensus between different distance metrics, based on which a single PCA will discover the similarity structure between the data points.

Figure 7.5: Evacuation time distributions for the (a) lattice gas, (b) social force, and (c) RVO2 models for $N = 50, 100, 150, 200, 300, 500, 1000$ agents. For each model and N , the distribution is built from 100 simulations.

7.3.1 Evacuation time

As Fig. 7.5 shows, the evacuation time distributions for all models have a uniform sub-distribution in the middle of the evacuation, which signals congestion at the exit, once we have $N > 50$ agents. However, in the RVO2 model this uniform sub-distribution is flanked by two small peaks. These two small peaks suggest a higher evacuation rate *just before* congestion sets in and *just before* congestion dissipates in the RVO2 model. This is not observed in real crowds, and hence is clearly an artifact of the model.

Comparing the three models, we find that the social force model has the longest evacuation times of up to five minutes. More importantly, careful inspection reveals fluid-like crowd dynamics around the exit in the RVO2 and lattice gas models. These can be seen in the videos http://www.youtube.com/watch?v=P64p3n1H_P4 and <http://www.youtube.com/watch?v=vJ0Nzi5Bykw>. In the RVO2 model, an obvious back flow can also be seen during the sudden gush of agents toward the exit. Due to the strong repulsion forces, we see a solid-like behavior in the social force model when the exit becomes congested (see video http://www.youtube.com/watch?v=30t7m959_yo).

In addition, the room was divided into six zones, as shown in Fig. 7.3,

Figure 7.6: Zoned evacuation time distributions for the (a) lattice gas, (b) social force, and (c) RVO2 models for $N = 1000$ agents in six zones, with zone 1 closest to the exit and zone 6 furthest away (see Figure 7.3). For each model, the distribution is averaged over 100 simulations.

based on the distance to the exit, to determine how strongly the crowds mix in the three models. The distributions of zoned evacuation times are shown in Fig. 7.6. These tell us that mixing is strong in the lattice gas and RVO2 models, but weak in the social force model. This mixing dynamics can be seen more clearly from the videos <http://www.youtube.com/watch?v=qeoJotgEUxk> (lattice gas), <http://www.youtube.com/watch?v=uZpd5L0DYZs> (RVO2), and http://www.youtube.com/watch?v=wEB6Ya0o_yw (social force).

From the movies of the lattice gas and RVO2 simulations, we see that agents from the nearer zones get pushed to the side once there is congestion at the exit. In these two models, agents prefer to keep moving when it is possible for them to do so. Agents in the social force model behave differently: once the exit becomes congested, they will become nearly stationary and wait for their turn to go through the exit.

Figure 7.7: Passage density maps of $N = 1000$ agents averaged over 100 simulations for the (a) lattice gas, (b) social force, and (c) RVO2 models. In these color maps, the exit is located at the bottom.

7.3.2 Passage density

The passage density heat map of agent locations averaged over 100 simulations provides information about the spatial-temporal trace of agents during their evacuation. This is crucial when analyzing the spatial structure of congestions. In Fig. 7.7a, we see a channel leading straight through the exit. This channel acts as an attractor, because once an agent gets onto this, it will be forced statistically towards the exit ($P_{t,y} = 1.00$). Though not as pronounced as for the lattice gas model, two left-right symmetric channels also formed in the RVO2 model (Fig. 7.7c). These two channels point roughly from the centers of mass of the left and right halves of the room towards the exit. No such artificial channels can be seen in the passage density map of the social force model (Fig. 7.7b).

Comparing the three models, we also see that the RVO2 model (Fig. 7.7c) has the most compact passage density, whereas the social force model (Fig. 7.7b) has the least compact passage density. In particular, in the region just in front of the exit, the RVO2 model gives rise to very high passage density. This high passage density cannot be attributed to agents stepping over a cell once or twice, as in the lattice gas and social force models. In the RVO2 model, this high passage density is produced by repeated visits to the same cell by individual agents that are always moving.

Figure 7.8: Distributions of total distance traveled for the (a) lattice gas, (b) social force, and (c) RVO2 models. For each model and N , the distribution is averaged over 100 simulations.

7.3.3 Total distance traveled distribution and inconvenience

To measure the difficulty (inconvenience) for an agent to squeeze through the crowd during evacuation, the *total distance traveled* and as well as the *inconvenience* (defined in Eq. 7.7) was measured. As we can see in Fig. 7.8b and Fig. 7.8c, the distributions of total distance traveled are very similar for the social force and RVO2 models. The lattice gas model, on the other hand, has very different distributions of total distance traveled (Fig. 7.8a). This is because space is continuous in the two former models, but discrete in the lattice gas model.

However, when we compare the distributions of inconveniences in Fig. 7.9, we find the lattice gas model is more qualitatively similar to the RVO2 model, though the typical inconveniences in the lattice gas model are larger than the those in the RVO2 model. For different N , the distribution peaks around the same inconvenience for both models. In contrast, the peak position of the inconvenience distribution of the social force model depends strongly on the number of agents in the simulation.

Figure 7.9: Distributions of inconveniences for the (a) lattice gas, (b) social force, and (c) RVO2 models. For each model and N , the distribution is averaged over 100 simulations.

7.3.4 Evacuation rate

The flow rate of agents at the exit was also measured. As shown in Fig. 7.10a, the flow rate of the lattice gas model saturates at around 11.8 persons per second due to the finite sizes of the agents and the exit. The social force model has the lowest flow rate of 2.2 people per second (Fig. 7.10b), because agents slow down dramatically at the end as a result of strong repulsive forces between agents. From Fig. 7.10c), we see the same peaks before and after the congestion period for the RVO2 model, a consequence of the liquid-like dynamics of RVO2 crowds.

In an attempt to compare simulation results with real data a number of videos showing real-life evacuations [188, 189] during the Sichuan Earthquake on May 12, 2008 were analysed. Two such videos, <http://www.youtube.com/watch?v=-y38ebiAnQw> and http://www.youtube.com/watch?v=e1yapP3z_L4 were analysed. Because the geometry of the scene and the viewing angles are not known, the flow rates could only be measured on a second-by-second basis from these two videos (Fig. 7.11), for comparison against what was measured in the simulations. Unlike the computational study, where the calculations could be averaged over 100

Figure 7.10: Flow rates at the exit as functions of time for the (a) lattice gas, (b) social force, and (c) RVO2 models. For each model and N , the flow rate is averaged over 100 simulations.

simulations to get smoothly varying flow rates, the flow rates measured from the videos are noisy. However, sensible comparisons can still be made. In <http://www.youtube.com/watch?v=-y38ebiAnQw>, the crowd is shown to evacuate through the check-out counters of a Walmart supermarket. The flow rate was high, but there was no congestion, because of the wide exit. This flow rate cannot be compared against the simulations of evacuation through a narrow exit.

In http://www.youtube.com/watch?v=e1yapP3z_L4, students were evacuating their classroom through a narrow door. From the video, we see that the width of the door allows no more than three school children to exit simultaneously. Therefore, the exit dimensions seem to correspond to the exit dimensions use in the simulations. From Fig. 7.11, we see that the average flow rate in the congestion phase is about 3 persons per second. This is far below the congestion flow rate in the lattice gas model, and very close to the congestion flow rate of the social force model. Based on this observable, we find that the social force model agrees best with the real data. This small test also tells us that the three models can be distinguished through their congestion flow rates, if other observables could not be measured.

Figure 7.11: The flow rates for evacuation during the May 2008 Sichuan Earthquake: (top) through the check-out counters of a Walmart supermarket in Chengdu, China (<http://www.youtube.com/watch?v=-y38ebiAnQw>); and (bottom) through the narrow exit of a classroom in an elementary school in Sichuan, China (http://www.youtube.com/watch?v=e1yapP3z_L4).

7.3.5 DISTATIS comparison

Finally, we make use of the DISTATIS method explained in Section 7.2.2 to compare the three models with all six different measurements. Fig. 7.12 shows the different measurements of the three models projected onto their first two principal components. The barycenters for the three models are also shown. From Fig. 7.12 we see that based on the total distance traveled (green), the social force and RVO2 models have similar distributions, which are both very different from that of the lattice gas model. On the other hand, the lattice gas and RVO2 models have similar distributions that are significantly different from that of the social force model.

More importantly, we find the clustering of some observables: evacuation time, passage density, and flow rate all tells us that the lattice gas model is similar to the RVO2 model, and different from the social force model. We also find the inconvenience and zoned evacuation time provide some indication that the three models are qualitatively different. In fact, the DISTATIS method indicates that the zoned evacuation time is one key metric that clearly distinguishes the three models. Therefore, any real world experiments should attempt to measure this as it is this observable maximally discriminates between the three models and therefore it is reasonable to expect it will also discriminate between real-world data and the models.

7.4 Summary

Motion planning systems that determine how an agent moves from one location to another form a fundamental part of almost all agent based crowd simulations. However, despite the number of models that exist, a standard methodology for quantitative

Figure 7.12: The DISTATIS analysis of the three models with $N = 1000$ agents. In this plot, τ is the quality of compromise, and λ is the eigenvalue.

comparison of existing motion planning systems is missing. This also implies that a standard method for quantitative validation against real world data does not exist. In this chapter, this crucial problem was addressed through the introduction of a methodology for quantitative comparison of motion planning systems used in egress simulation.

More specifically, emergency evacuation from a rectangular room with a single narrow exit was simulated using three models: (1) lattice gas, (2) social force, and (3) RVO2, with different starting number of agents in the room. From these simulations, six observables were measured: (1) evacuation time, (2) zoned evacuation time, (3) passage density, (4) total distance traveled, (5) inconvenience and (6) flow rate. Besides qualitative comparison of the various distributions obtained, the three models were also compared quantitatively using the DISTATIS method. Comparing the evacuation time, passage density, and flow rate distributions, we find that the lattice gas and RVO2 models are similar, both of which are very different from the social force model. On the other hand, if we compare the distribution of total distance traveled, the social force and RVO2 models are more similar to each other, and very different from the lattice gas model. When the simulated congestion flow rates of the three models were compared against the congestion flow rates of school children evacuating from their classroom during the May 2008 Sichuan Earthquake, we find that the social force model agrees best with the egress data. Finally, from an analysis

of the DISTATIS the zoned evacuation time can be identified as the one observable metric that can best discriminate between these models, and also between models and real-world data. This could be invaluable in verifying new models and even existing models when new and more detailed video data is available.

Chapter 8

Conclusion and Future Work

Crowd egress from buildings is a complex process. Understanding and analyzing the behavior of crowds during egress is crucial in order to prevent disasters like the Love Parade and the Iroquois Theater from being repeated. Modelling and simulation is a primary tool for doing this. Modern agent based models, in particular, can have complex behavior models that enable the simulation and analysis of egress in much higher detail than was ever possible. However, computational models are still unable to simulate the entire process of egress from the sound of the fire alarm to the last person evacuating safely. This thesis took a holistic approach to understanding and modelling the egress process. Literature was examined from all related fields, including psychology, civil engineering, sociology, etc. Along with this understanding, incident reports and analyses were used to clearly identify the core components of the egress process and discover the steps in that process that are critical to improving safety. Chapter 3 describes the details of the egress process gained from this research. With this fundamental understanding it was then possible to evaluate the existing computational models commonly used when modelling human egress. This comparison lead to the identification of key research questions in which either new modelling approach or new understanding was necessary. The contributions of this thesis was to develop these areas and thereby improve the state of the art in the computational modelling of egress. In what follows the identified research gaps are summarised.

The commonly held conception of emergency egress is that during an emergency, people panic, stampede and head for the closest exit and, in the process, injure and kill others. However, the consensus among crowd behavior experts who've studied past incidents is that people do not *panic* or behave irrationally in a crowd (see chapter 2). People generally tend to behave rationally within the bounds of their knowledge. This has two important implications for egress modelling: Firstly, modelling rational, thinking human beings would be necessary for realistic egress simulation; secondly, the complete process of how people gain information from the environment and react to it will have to be considered to model behavior during egress. Simply considering agents that rush straight for the exit at the sound of the

fire alarm (which is generally considered the start of an evacuation simulation) would be insufficient to accurately simulate egress.

A study of incident reports, psychology and sociology studies and experiments in psychology (see Chapter 2) revealed a great deal about human behavior during egress. Occupants rarely have a complete understanding of the situation when an emergency starts. Most occupants do not have complete knowledge of the spatial layout of the building or the best routes for escape. In fact, many studies report the significant amount of time that is spent by occupants, engaging in pre-evacuation behavior of first disbelieving the first cues that are observed and subsequently investigating for more information. This search for information is further complicated by the effect of stress and time constraints which limit the information processing capacity of the brain. The human brain compensates for this by chunking together similar information and only processing what is determined to be the most significant and important of these cues. Once the need to evacuate is determined, occupants try to search for familiar and reassuring surroundings by trying to find their family or friends, i.e. their primary group. Once found, they try to exit by the shortest route *that they know of*. This is further complicated by the fact that most occupants do not know the complete layout of the building. This leads to inefficient routes and exploration behavior. Eventually, through following others and, generally, with the help of trained personnel occupants escape from the building.

In general, the study of human behavior literature seemed to indicate that most of the process of human behavior during egress is well understood. Several computational models of varying complexity have been developed over the years to simulate this egress behavior. In order to study these models, they were categorized according to the approach used for modelling behavior and the level of detail in which the crowd is treated. Agent based models, with their approach to considering individual agents and their interactions, were found to come closest to simulating the complex process of egress behavior. Several highly capable models exist that are able to simulate the effect of emotion, leadership behavior and the impact of fire and smoke on mobility and visibility. However, existing models do not even attempt at simulating pre-evacuation behavior or the effects of human cognitive limits on other aspects of human behavior like spatial knowledge utilization and movement. In the few cases, that this is attempted, it is observed that the models were too simple to help in actually analyzing the effects of these information processing limits on egress.

As a first step to modelling and analyzing the effect of human cognitive limits on egress behavior, the four key components of the egress process were identified: Perception, Event Identification, Spatial Knowledge Utilization and Navigation. There are several issues that arise in modelling and analyzing the effect of cognitive limits on each of these components. In some cases, like *event identification* and *perception*, there is lot of existing literature in psychology and cognitive science

on the effects of these cognitive limits; however, they have never been considered or modeled in existing computational models. In these cases, Chapters 4 and 5 proposed novel models for these and demonstrated the usefulness of considering these in egress modelling. In other cases, like *spatial knowledge modelling*, a study of the literature revealed that we do not yet have a sufficient understanding of exploration behavior in the presence of lack of spatial information and the role of memory in wayfinding. Thus, in Chapter 6, a scalable game based methodology was developed to learn more about the role of memory and the strategies used by people in indoor wayfinding. Finally, in the case of *navigation* modelling, there exist a plethora of models taking different approaches to navigation; however, a standard methodology for comparing and analyzing these different models has not been developed. This makes it difficult to determine the effectiveness or usefulness of a new model even if it is developed. This research problem was addressed in Chapter 7. In the following sections of this final chapter, the work done in each of these chapters is briefly reviewed with their limitations and future possibilities.

8.1 Modelling Human Cognitive Limits in Perception

The first component of simulating human behavior in egress is the model of how the environment is perceived. Perception is a fundamental part of any agent based model of crowds. Existing crowd simulation models simply consider perception as agents perceiving all objects within a certain circular (or at times an elliptic) sensing range. However, as mentioned earlier, limitations of human short term memory mean that a person is able to process no more than nine percepts at any given time. In fact, more recently, it's been shown that this limit is probably closer to five than nine. One of the ways in which the human mind counters this limitation is by chunking together similar information thus making better use of its limited capacity. We believe that modelling this is necessary to get more realistic simulations of human behavior.

In Chapter 4, a model for perception that takes into consideration this limited information processing capacity and chunking of information was introduced. The novel idea was in considering sensing as a process of gathering *information* from the environment rather than simply visual percepts for collision avoidance. Its working and usefulness was demonstrated by enhancing the popular Reciprocal Velocity Obstacle Based motion planning system with this Information-processing Based Perception (IBP) System. The chunking of information was modeled through a global dynamic clustering that is performed once per time-step. Through this efficient group based perception system, the way in which people avoid groups of people coming towards them was modeled. Multiple layers of clustering were used to ensure that simulated pedestrians did not move unnaturally when avoiding big groups. Comparison against

real world experiments and analysis of simulation results in several standard scenarios showed the modelling perception as a process of gathering information helped produce more realistic results. Simulations were also used to show how a navigation system using IBP could work in a system where the limited information processing capacity of humans was considered. In principle, IBP could be used to model all forms of perception of an agent while still taking into consideration the cognitive limits of short term memory.

One of the primary aims of the development of IBP was to combine it with a cue perception system while still taking into consideration the limited information processing capacity of human beings. A similar filtering system could be used for determining the information value of cue percepts and other percepts and only a limited number of the most important percepts would be processed by the agents. Theoretically, this is likely to produce more realistic behavior and also the bounded rationality that is discussed in literature on behavior during emergencies. However, there are several factors that are very difficult to consider in doing this, the most pertinent of which would be the difficulty in assigning an information value to different percepts.

Ideally, the proposed model would have been tested against more video evidence of crowds. However, this data is not easy to obtain and validate against. Another limitation of the work, due to time constraints, is that most of the simulations performed were rather simplistic compared to the multi floor building simulations in Chapter 5.

Perhaps, one of the most interesting products of this research is the basis it provides for modelling cognitive limits in computational models of humans. The idea of grouping together similar information during perception to model scientifically observed *chunking* behavior provides us with better insight into the implications of the extensive work done in cognitive science. In the context of egress simulation, it provides planners with a method for analyzing and determining how information can be passed to evacuees during emergencies where our ability to perceive information from the environment becomes even more constrained due to stress and time constraints. In the following chapter, a computational model for how occupants process perceived information during egress was developed.

8.2 A Novel Model for Event Identification

In Chapter 5, a method for modelling the processing of cues and identifying events was described. It was demonstrated how this model could be used to model pre-evacuation behavior and the effect of different kinds of fire alarms and other cues. Experiments demonstrated the capabilities of the model and, more importantly, identified the necessity of modelling of pre-evacuation behavior. For example, they helped identify

the critical number of occupants that need to be trained if they were randomly distributed over a two storey environment. The impact of milling behavior on evacuation dynamics was also shown. The model was also used to show that strategically placed, trained personnel may be just as effective as an expensive fire alarm in improving survival rate.

This bucket based model of event identification opens up several possibilities for future work. Some of these could be done through modifications to the experiments conducted and were suggested in the chapter itself. For example, the model could be used for identifying strategies for distributing trained personnel. A preliminary exploration into the usefulness of strategic distribution was in fact presented as well. Different strategies of exploration for trained personnel could be explored to ensure maximum survival or the amount of time they should spend finding others before evacuating themselves could be determined if the building structure is known. The experiment on milling behavior could be used during building construction to ensure optimality of milling/ gathering positions. In case of budget constraints, the optimal distribution of expensive and unambiguous fire alarm systems and cheaper more ambiguous alarms could be determined while ensuring maximum effectiveness.

The major limitation of this study was the lack of comparison against real world data. Considering the nature of the scenarios under consideration, there is little that can be done to counter this. However, with the advent and popularity of wearable electronic devices and the array of sensors on modern smartphones there is a possibility that this data might be available in the future. At present, the best possible solution would be to use a game based methodology like the one proposed in Chapter 6.

As discussed in Section 8.1, another avenue of future work would be to combine the cue perception and event identification system with the Information Based Perception system and to analyze the effectiveness of this combined model. This may also help identify further shortcomings in both these models. More importantly, it would open doors to using modelling and simulation to analyze how information can be passed to occupants to improve egress efficiency; for example, more complex fire alarm systems and interactive sign boards could be tested and improved cheaply before a huge investment is made. This kind of analysis, however, is difficult to do with our current understanding of the effect of partial spatial knowledge on indoor wayfinding. This was further explored in the following chapter.

8.3 A Game Based Analysis of Indoor Wayfinding Behavior

As mentioned earlier, one of the difficulties in modelling the effect of the limits of human information processing capacity on spatial knowledge utilization is our limited

understanding of the phenomenon. Existing models of simulation take the simple approach of assuming that all evacuees have complete knowledge of the environment. When partial knowledge is modeled, a simple approach of having perfect memory and remembering all rooms perceived once is modeled. Also, exploration behavior has rarely been explored in any detail in existing crowd simulation models. A review of existing literature in cognitive psychology revealed that a lot of work had been done on human spatial memory and exploration of outdoor environments. However, exploration and way-finding in multi-floor indoor environments is still not very well understood. It was also discovered that while interesting experiments have been developed to study indoor wayfinding, these experiments are difficult to scale and tend to analyze only a few participants. This makes it difficult to do more detailed analysis.

To explore and understand the role of memory in indoor wayfinding, a game was developed in Minecraft which required players to explore a multi-floor building and complete certain tasks within certain time constraints. Using Minecraft had several advantages including the popularity of the game, the ease of changing the environment and, most importantly, the possibility of tracking all the actions of players. This data was then analysed using some agent based models to determine strategies of exploration and spatial memory limitations. It was found that, during exploration, people prefer a floor first strategy and try to avoid making decisions to change direction unless they are forced to. A novel Markovian analysis was developed to analyze the size of memory used during exploration by the players and revealed that 6 – 8 step memory is sufficient to produce human like exploration in agents.

One of the limitations of the analysis discussed in this chapter was the amount of data that was available for analysis especially for the Markov analysis. The limited number of participants available also meant that it was difficult to test hypothesis on the task completion phase by creating and comparing the players movement in an alternate environment.

As discussed in Section 6.3.2, one of the big advantages of using a popular game like Minecraft with over 35 million copies sold, is the ubiquity of the game. To make use of this ubiquity, the game has been hosted on a Linux Server provided by Amazon Web Services and data will be collected from this server over the next year. First a simple web page (vaisaghvt.com/minecraft-experiment/) was created where interested players can submit a request to play by using a registered Minecraft account. This request is sent to the Amazon Server. Python's event driven network engine *Twisted* [190] was used to create a simple application that listens to these requests at a specified port and starts a white-listed Minecraft Server to which only this particular user can connect. Once a request is successful, the player receives the IP address of the machine which he proceeds to use to connect to the server and play the game. The Minecraft Server is shut down once the player exits the game.

Due to the cost of running a virtual server with 3 GB memory, we have only one server and thus only one player is able to connect at any given time. With about 1-3 players a day, we hope to collect data from 500-1000 players over the next year. This much greater sample will hopefully reveal more fundamental and interesting aspects of human exploration behavior than was possible to analyze with the limited dataset currently available. The setup of the experiment also makes it relatively easy to scale to multiple servers if they are available. Running different experiments just involves replacing a single folder with a different setup. With enough data, this can be used to do A/B testing to test different hypothesis. The game could itself be changed in several ways to study other aspects of indoor navigation. For example, by removing the limitation that only one player can connect to the server at any given time, multi-occupant evacuation from the indoor environment and the communication strategies used during egress could be studied.

A limitation of the scaled experiment being hosted online is the dependence on the integrity of the players. It is near impossible to check for all kind of cheating. Also it is more difficult to do traditional methods like asking players to speak their thoughts out loud to determine their thought process. A more complicated incentive system which pays players some amount based on their achievements may have to be developed to ensure participation.

In spite of these limitations, even on a local server, the game based methodology was able to gather a larger participant base than existing desktop virtual reality experiments. This proved invaluable in using the Markov Agent Based analysis which provided novel insights into the role of short term memory in indoor wayfinding. This, in turn, enables the more accurate modeling of spatial knowledge utilization and indoor exploration in agent based simulations of egress.

8.4 A Comparison of Existing Motion Planning Systems

The fourth and final component of the egress process is the navigation system. This part of the model simulates how the agent moves from a particular source location to it's destination while avoiding collisions with other agents. The navigation system can generally be divided into two parts: a higher level path planning system and a lower level motion planning system. The path planning system determines a high level path towards a selected goal, usually in terms of some way-points. The motion planning system which is sometimes also called a collision avoidance system helps the agent select a collision free velocity while moving towards the next way-point. There has been a lot of research on this and there are several different motion planning systems that exist today. A problem that lies in determining or developing the most realistic motion planning system among these is the lack of a method for the systematic com-

parison of these different motion planning systems. In Chapter 7, a methodology for comparing different models of crowds was introduced and the DISTATIS method was used to identify zoned evacuation time as a key metric that can be used in differentiating the different models. Video footage from the Sichuan Earthquake in China was then used for determining which model produced the most *realistic* behavior.

The simulations in this research were done only in a simple square room with a single exit. Similar analysis in more complex layouts may help reveal more about the models being compared. Also, only three of the most popular models i.e. RVO2, Social Force and Lattice Gas models were compared in this study. More models will probably need to be compared to increase confidence in the notion of zoned evacuation time being the key metric for differentiation between models.

Due to the limited amount of video footage available and the difficulty in extracting meaningful statistics from these videos, only the flow rate could be determined for the real world data and compared against the models. Even though zoned evacuation time was identified as the key metric, it was impossible to measure this from the limited video data available.

Other metrics that weren't considered in the chapter could also be used for comparing the models and a similar DISTATIS method can be used for discovering a new key metric. We do not hypothesize that the methodology proposed is the final arbiter on deciding which motion planning system to use. When actually using a motion planning system other considerations will have to be made as well like the purpose of the simulation, the computational resources required, the ease of development and its compatibility with the other systems in the simulation. However, the methodology developed does provide a novel and systematic method for comparison of different models and validating new models of navigation that are developed.

8.5 Future Work

A problem repeatedly encountered in the agent based modeling of crowd egress is validation. One of the primary reasons for this is the scarcity of the right kind of data and the difficulty in collecting this data. Two recent technological developments have opened possibilities for the development of novel methodologies to tackle this problem: commercial Virtual Reality (VR) headsets and the ubiquity of sensors.

Immersive Virtual Reality using VR headsets provides a way in which the immersiveness of traditional real world experiments can be combined with the experimental control and data-tracking capability of serious game methodologies of the kind used in Chapter 6. Until recently, however, development of VR environments of similar complexity to desktop virtual environments was difficult due to the limitations in the tools available. However, the advent of commercial VR headsets like Oculus Rift and Sony's Morpheus Project brings a lot of opportunities. Cheaper equipment

would mean much more ubiquity and the possibility of the development of scalable games of the kind developed in Chapter 6. More importantly, the APIs available on these equipment mean that more complicated experiments can be developed with better tracking capabilities than was possible before. For example, VR based games could be used for validating and calibrating the proposed Information Based Perception system through gaze tracking; or different aspects of the pre-evacuation behavior could be tested without the limitations of real world experiments.

The ubiquity of sensors in the form of smartphones, wearable electronics and other smart devices that form the internet of things provides possibilities for novel methods for learning more about crowds and validating models. For example, the zonal evacuation time that was identified as a key metric for differentiating between models and comparing them against real world data could be measured in simple real world evacuation experiments with participants wearing tags that help determine the evacuation rate for different zones. Furthermore, wearable electronics could be used to supplement the limited perception capacity of human beings by providing relevant information in evacuations or high density situations. Crowd simulations could be used to guide the development of these kind of devices that could prove invaluable in preventing disasters due to crowding [191]. Symbiotic data-driven simulations [192] that adapt according to the data from sensors and feed back information to the system that facilitates efficient egress could be a real possibility soon.

8.6 Conclusion

Over the course of this thesis, methods that enable the modelling and analysis of the effect of human cognitive limits on the egress process has been developed. This chapter has reviewed the major contributions of this thesis with brief discussions of the several possible extensions and new areas of research that have opened up as a result of this thesis. The work done in this thesis not only advanced the state of the art in the computational modelling of egress but has also helped improve our understanding of the different components of human behavior during emergency egress. This could prove invaluable in saving lives by helping planners understand, manage and control dangerous situations that develop with crowding.

Appendix A

Comparison of Information Based Perception against Controlled Experiments

Chapter 4 introduce an information based perception system and demonstrated its effectiveness in helping existing motion planning systems produce more realistic behavior. In the chapter, results from simulations were compared against the controlled experiments conducted by Hu et al. [1, 129]. In this section we present the results of two additional metrics — the speed as a function of time and the trajectory length — and compare the results produced against traditional RVO2, social force and the results produced in the experiment.

Regarding trajectory length, Hu’s most significant finding was that, typically, the oncoming agent would travel a longer distance in order to avoid the oncoming group as shown in Table A.1. All three models produced similar dynamics as shown in Table A.2.

Figure A.1 shows the development of speed against time for the three models that were considered. Hu et al observed that one of the problems with existing motion planning systems was that there was a significant slowing of agent 2 when it gets closer to the other agents; whereas in the real world results all three participants maintained their speed. When RVO2 is used with IBP there is a dip in speed of

Trajectory Length Calculation from Experiments			
	Category 1 (55%)	Category 2 (40%)	Category 3 (5%)
Agent 1	10.0542	10.12006	10.15071
Agent 2	10.238	10.27922	10.15928
Agent 3	10.1126	10.0789	10.18924

Table A.1: Trajectory length according to experiments conducted by Nan [1]. Category 1 refers to the movement of Agent 2 to the left of the group, category 2 to the movement of Agent 2 to the right of the group and category 3 to the movement of agent 2 through the middle of the oncoming group.

	Experiment (95%)	RVO2+IBP	RVO2	Social Force
Agent 1	10.087	9.925	9.925	10.003
Agent 2	10.259	12.132	11.833	11.831
Agent 3	10.096	9.925	9.925	10.003

Table A.2: This table shows the trajectory length obtained for the simulation of the experimental scenario in [1]. The first column lists the average of category 1 and 2 which together comprised 95% of the real world experiment data of [1].

(a) Agent 2

(b) Agent 1

(c) Agent 3

Figure A.1: Hu et al. observed that real world participants reached a stable average speed of 1.5m/s^{-1} and maintained that speed for the duration of the experiment run. This figure shows the results produced by RVO2, Social Force and IBP-RVO2. While Social Force seems to produce results that agrees with the experiment, Agent 2 seems to slow down considerably when traditional RVO2 is used. However, when used with IBP, there is only an instantaneous drop in velocity (for less than a second) as in Social Force.

participant 2 but this lasts for less than a second as shown in Fig A.1. For the rest of the simulation, the agent maintains its speed. As seen in the figure, this isn't the case with RVO2 with a traditional sensor range.

Appendix B

Validity Check for Markov analysis

Chapter 6 introduced a novel Markovian analysis to analyse the use of memory in exploration. The effectiveness and validity of the Markovian analysis is, however, dependent on the amount of data available. The following experiments demonstrate the validity of the results obtained from the analysis.

B.1 Effect of Dataset Size

In order to check whether the peak in the hop and coverage data is a function of the dataset size, we hypothesize that if the peak does not change on doubling the dataset size then the pattern that is seen is not an artifact of the dataset size. The same graph for $N = 22$ (i.e. taking only 22 participants at random) and $N = 44$ was plotted and checked for a shift in the peak to a higher value. As shown in Figure B.1 there is clearly no shift. The peak value still remains at 7-8 and starts dropping at 9.

B.2 Decision Base Size at Decision Points

The Markov data calculation (explained in Section 6.4.2) is based on aggregating the actions of the players at each decision point. The quality of the Markov data probability distribution calculated is thus based on the amount of data available to make predictions. In order to calculate an m^{th} order probability of the next step, the calculations are done using the number of times the preceding sequence of m steps occurred and the choices made at that point by players in the dataset. Thus, as the frequency of said path decreases, there is lesser data for the probability calculation. The m^{th} order Markov calculation becomes progressively less reliable and less different from random noise as m increases due to the limited size of the dataset.

To estimate the maximum value of m for which m^{th} order calculations can be done reliably we calculate the *decision base size* of the calculations.

To calculate the decision base size for order m , we do the following: First, 30,000 paths of 300 hops are generated for an m^{th} order Markov agent based

(a) Hop and coverage for 44 players

(b) Hop and coverage for 22 players

Figure B.1: Hop and coverage graphs for different size of datasets seems to indicate that the peak is not dependent on data size.

Figure B.2: Comparison of the decision base sizes.

on the dataset of size N . For each node in each path generated, the number of players that have actually made a decision at that point is counted. We divide this value by

the number of decisions that are possible at that point. So a single decimal value is obtained for each node of each path generated. The average of this over all nodes a path gives the *path specific decision base*. The average of these path specific decision bases gives the decision base for the order m and dataset of size N . The result of this calculation is shown in Figure B.2.

The figure, firstly and most importantly, shows the usefulness of having a larger dataset in the Markov data calculation. Also, beyond the 9th order a decision base of size less than 2 is available indicating there are less than twice as many actions as there are choices on average at each decision point. This indicates that beyond the 9th order the calculations are unreliable.

References

- [1] N. Hu, “Spatial-Temporal Patterns and Pedestrian Simulation,” Ph.D. dissertation, Nanyang Technological University, May 2014.
- [2] D. Helbing and P. Mukerji, “Crowd disasters as systemic failures: analysis of the love parade disaster,” *EPJ Data Science*, vol. 1, no. 1, pp. 1–40, 2012. [Online]. Available: <http://dx.doi.org/10.1140/epjds7>
- [3] G. K. Still, “Crowd Dynamics,” Ph.D. dissertation, University of Warwick, Department of Mathematics, Aug. 2000.
- [4] S. Zhou, D. Chen, W. Cai, L. Luo, M. Y. H. Low, F. Tian, V. S.-H. Tay, D. W. S. Ong, and B. D. Hamilton, “Crowd modelling and simulation technologies,” *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, vol. 20, no. 4, pp. 20:1–20:35, Nov. 2010. [Online]. Available: <http://doi.acm.org/10.1145/1842722.1842725>
- [5] S. Gwynne, E. Galea, M. Owen, P. Lawrence, and L. Filippidis, “A review of the methodologies used in the computer simulation of evacuation from the built environment,” *Building and Environment*, vol. 34, no. 6, pp. 741 – 749, 1999. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0360132398000572>
- [6] S. Regelous and K. Mannion. (2011) Massive Software – Simulating Life. [Online]. Available: <http://www.massivesoftware.com/>
- [7] C. W. Reynolds, “Flocks, herds and schools: A distributed behavioral model,” *Computer Graphics*, vol. 21, no. 4, pp. 25–34, Jul. 1987.
- [8] J. Snape, S. J. Guy, M. C. Lin, and D. Manocha, “Local and global planning for collision-free navigation in video games,” *Planning in Games Workshop*, p. 7, 2013.
- [9] R. E. Ensemble Studios, Big Huge Games, “Age of empires,” April 2013. [Online]. Available: <http://ageofempiresonline.com/en/>
- [10] D. Helbing and P. Molnar, “Social Force Model for Pedestrian Dynamics,” *Physical Review E*, vol. cond-mat.stat-mech, pp. 4282–4286, May 1995.

- [11] V. Viswanathan and M. Lees, “Modelling and analyzing the human cognitive limits for perception in crowd simulation,” in *Transactions on Computational Science XVI*, ser. Lecture Notes in Computer Science, M. Gavrilova and C. Tan, Eds. Springer Berlin Heidelberg, 2012, vol. 7380, pp. 55–76. [Online]. Available: http://dx.doi.org/10.1007/978-3-642-32663-9_4
- [12] S. J. Guy, J. Chhugani, S. Curtis, P. Dubey, M. C. Lin, and D. Manocha, “PLEdestrians: a least-effort approach to crowd simulation,” in *SCA '10: Proceedings of the 2010 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, University of North Carolina. Eurographics Association, Jul. 2010.
- [13] H. Klüpfel, M. Schreckenberg, and T. Meyer-König, “Models for crowd movement and egress simulation,” pp. 357–372, 2005.
- [14] R. D. Peacock, E. D. Kuligowski, and J. D. Averill, *Pedestrian And Evacuation Dynamics*. Springer, 2011. [Online]. Available: <http://link.springer.com/book/10.1007/978-1-4419-9725-8/page/1>
- [15] A. Mordvintsev, V. Krzhizhanovskaya, M. Lees, and P. Sloot, “Simulation of city evacuation coupled to flood dynamics,” in *Pedestrian and Evacuation Dynamics 2012*, U. Weidmann, U. Kirsch, and M. Schreckenberg, Eds. Springer International Publishing, 2014, pp. 485–499. [Online]. Available: http://dx.doi.org/10.1007/978-3-319-02447-9_40
- [16] W. B. Arthur, “Complexity in economic and financial markets: Behind the physical institutions and technologies of the marketplace lie the beliefs and expectations of real human beings,” *Complexity*, vol. 1, no. 1, pp. 20–25, 1995.
- [17] M. Kobes, I. Helsloot, B. de Vries, and J. G. Post, “Building safety and human behaviour in fire A literature review,” *Fire Safety Journal*, vol. 45, no. 1, pp. 1–11, Dec. 2009.
- [18] A. Schadschneider, W. Klingsch, H. Klüpfel, T. Kretz, C. Rogsch, and A. Seyfried, “Evacuation dynamics: Empirical results, modeling and applications,” in *Encyclopedia of complexity and systems science*. Springer, 2009, pp. 3142–3176.
- [19] S. Reicher, “The Psychology of Crowd Dynamics,” in *Blackwell Handbook of Social Psychology: Group Processes*, M. A. Hogg and R. S. Tinsdale, Eds. Blackwell Publishers Ltd, Jan. 2008, pp. 182–208.
- [20] M. R. Torres, “Every Man for Himself? Testing Multiple Conceptual Approaches of Emergency Egress on Building Evacuation During a Fire,” Ph.D. dissertation, University of Delaware, Apr. 2010.

- [21] R. L. Paulsen, "Human behavior and fires: An introduction," *Fire Technology*, vol. 20, pp. 15–27, May 1984.
- [22] J. D. Sime, "Affiliative Behavior During Escape to Building Exits," *Journal of environmental psychology*, vol. 3, pp. 21–41, 1983.
- [23] B. E. Aguirre and G. Santos, "A Critical Review of Emergency Evacuation Simulation Models," in *Building Occupant Movement During Fire Emergencies*. Gaithersburg, Maryland: University of Delaware, 2004, pp. 1–54.
- [24] J. D. Sime, "Crowd psychology and engineering* 1," *Safety Science*, vol. 21, pp. 1–14, 1995.
- [25] G. A. Miller, "The magical number seven, plus or minus two: some limits on our capacity for processing information." *Psychological review*, 1956.
- [26] F. Ozel, "Time Pressure and Stress as a Factor During Emergency Egress," *Safety Science*, vol. 38, no. 2, pp. 95–107, 2001.
- [27] R. L. Davis, B. A. Therrien, and B. T. West, "Working memory, cues, and wayfinding in older women," *Journal of Applied Gerontology*, vol. 28, no. 6, pp. 743–767, 2009. [Online]. Available: <http://jag.sagepub.com/content/28/6/743.abstract>
- [28] D. Berry and R. Hanley. (2000, Jan.) Questions About Fire Safety After Deaths at Seton Hall . [Online]. Available: <http://www.nytimes.com/2000/01/21/nyregion/questions-about-fire-safety-after-deaths-at-seton-hall.html>
- [29] J. Tsai, N. Fridman, E. Bowring, M. Brown, S. Epstein, G. Kaminka, S. Marsella, A. Ogden, I. Rika, A. Sheel *et al.*, "Escapes: evacuation simulation with children, authorities, parents, emotions, and social comparison," in *The 10th International Conference on Autonomous Agents and Multiagent Systems-Volume 2*. International Foundation for Autonomous Agents and Multiagent Systems, 2011, pp. 457–464.
- [30] N. Pelechano and N. I. Badler, "Modelling Crowd and Trained Leader Behavior during Building Evacuation," *IEEE Computer Graphics and Applications*, vol. 26, no. 6, pp. 80–86, Dec. 2006.
- [31] T. T. Pires, "An approach for modelling human cognitive behavior in evacuation models," *Fire Safety Journal*, vol. 40, no. 2, pp. 177–189, Mar. 2005.
- [32] H. L. Klüpfel, "A Cellular Automaton Model for Crowd Movement and Egress Simulation," Ph.D. dissertation, Universität Duisburg-Essen, der Fakultät 4 – Naturwissenschaften, Jul. 2003.

- [33] F. A. Aveni, "The Not-So-Lonely Crowd: Friendship Groups in Collective Behavior," *Sociometry*, vol. 40, no. 1, pp. 96–99, Mar. 1997.
- [34] P. M. A. Slood, A. Schoneveld, and J. De Ronde, "Large-scale simulations of complex systems, part I: conceptual framework," *Working Papers*, 1997.
- [35] E. D. Kuligowski, "The Process of Human Behavior in Fire," National Institute of Standards and Technological, Tech. Rep. 1632, May 2009.
- [36] K. Andrée and B. Eriksson, "Cultural differences in an evacuation scenario," Ph.D. dissertation, Lund University, 2008.
- [37] A. Sandberg, *Unannounced evacuation of large retail stores*, ser. An evaluation of human behaviour and the computer model SIMULEX, R. Jonsson and T. J. Shields, Eds. Lund Institute of Technology, Lund University: Lund University, Oct. 1997.
- [38] N. Pelechano, K. O'Brien, B. Silvermann, and N. I. Badler, "Crowd Simulation Incorporating Agent Psychological Models, Roles and Communication," in *International Workshop on Crowd Simulation*. University of Pennsylvania, 2005, pp. 21–30.
- [39] H. Aydt, M. Lees, L. Luo, W. Cai, M. Y. H. Low, and S. K. Kadirvelen, "A Computational Model of Emotions for Agent-Based Crowds in Serious Games," in *Intelligent Agent Technologies*. Lyon, France: School of Computer Engineering, Nanyang Technological University, 2011, pp. 1–9.
- [40] X. Pan, "Computational modelling of Human and Social Behaviors For Emergency Egress Analysis," Ph.D. dissertation, Stanford University, The Department of Civil and Environmental Engineering, Stanford University, Jun. 2006.
- [41] R. d. S. França, M. d. G. B. Marietto, and M. B. Steinberger, "A Multi-agent Model for Panic Behavior in Crowds," in *Fourteenth Portuguese Conference on Artificial Intelligence*. Federal University of ABC, Santo André, São Paulo, Brazil, Oct. 2009, pp. 463–472.
- [42] C. Cocking and J. Drury, "DON'T PANIC! Altruistic behaviour in crowd emergencies," in *Altruism Conference*, Oct. 2005, pp. 1–11.
- [43] G. Proulx, "Occupant behaviour and evacuation," in *9th International Fire Protection Symposium*. Munich: Institute for Research in Construction, National Research Council Canada, May 2001, pp. 219–232.
- [44] G. Ramachandran, "Human behavior in fires: a review of research in the United Kingdom," *Fire Technology*, pp. 149–155, May 1990.

- [45] A. Mawson, "Understanding mass panic and other collective responses to threat and disaster," *Psychiatry: Interpersonal and Biological Processes*, pp. 95–113, 2005.
- [46] G. Proulx, "Playing with fire understanding human behavior in burning buildings," *ASHRAE Journal*, vol. 45, no. 7, pp. 33–35, Jul. 2003.
- [47] D. Tong and D. Canter, "The Decision to Evacuate a Study of the Motivations which Contribute to Evacuation in the Event of Fire Theoretical deficiencies," *Fire Safety Journal*, vol. 9, pp. 1–9, 1985.
- [48] T. L. Graham and D. J. Roberts, "Qualitative overview of some important factors affecting the egress of people in hotel fires," *International Journal of Hospitality Management*, vol. 19, pp. 79–87, Feb. 2000.
- [49] G. Proulx, "Evacuation Time and Movement in Apartment Buildings," *Fire Safety Journal* 24, vol. 24, pp. 229–246, 1995.
- [50] D. A. Purser, "Quantification of behaviour for engineering design standards and escape time calculations," *Safety Science*, vol. 38, pp. 157–182, 2001.
- [51] G. Proulx, "High-rise office egress the human factors," in *Symposium on High-Rise Buidling Egress*. New York City: Institute of Research in Construction, National Research Council, Canada, May 2007, pp. 1–6.
- [52] B. Latané and J. M. Darley, "Bystander "Apathy"," *American Scientist*, vol. 57, no. 2, pp. 1–12, Jul. 1969.
- [53] C. Cocking and J. Drury, "The Mass Psychology of Disasters and Emergency Evacuations: a Research Report and Implications for the Fire and Rescue Service," *Fire Safety : Technology and Management*, vol. 10, pp. 13–19, 2008.
- [54] E. Quarantelli, "The Nature and Conditions of Panic," *American Journal of Sociology*, vol. 60, no. 3, pp. 267–275, Nov. 1954.
- [55] J. Drury, C. Cocking, S. Reicher, A. Burton, D. Schofield, A. Hardwick, D. Graham, and P. A. Langston, "Cooperation Versus Competition in a Mass Emergency Evacuation: a New Laboratory Simulation and a New Theoretical Model," *Behavior Research Methods*, vol. 41, no. 3, pp. 957–970, Aug. 2009.
- [56] C. Cocking, J. Drury, and S. Reicher, "The Psychology of Crowd Behavior in Emergency Evacuations: Results From Two Interview Studies and Implications for the Fire & Rescue Service," *The Irish Journal of Psychology*, vol. 30, pp. 1–31, 2007.
- [57] B. Cornwell, "Bonded Fatalities: Relational and Ecological Dimensions of a Fire Evacuation," *The Sociological Quarterly*, vol. 44, no. 4, pp. 617–638, Oct. 2003.

- [58] J. Chertkoff, R. Kushigian, and M. A. McCool Jr, "Interdependent exiting: The effects of group size, time limit, and gender on the coordination of exiting," *Journal of environmental psychology*, vol. 16, pp. 109–121, 1996.
- [59] B. D. Jones, "Bounded Rationality," *Annual Review of Political Science*, vol. 2, pp. 297–321, 1999.
- [60] R. Fahy and G. Proulx, "Toward Creating a Database on Delay Times to Start Evacuation and Walking Speeds for Use in Evacuation modelling," in *2nd International Symposium on Human Behavior in Fire*. Boston, MA, USA: National Research Council Canada, 2010, pp. 175–183.
- [61] M. J. Kinsey, E. R. Galea, P. J. Lawrence, D. Blackshields, L. Hulse, R. Dat, and G. Sharp, "Modelling pedestrian escalator behaviour," *Pedestrian and Evacuation Dynamics 2008*, Dec. 2009.
- [62] S. M. Olenick and D. J. Carpenter, "An Updated International Survey of Computer Models for Fire and Smoke," *Journal of Fire Protection Engineering*, vol. 17, pp. 87–110, May 2003.
- [63] P. A. Langston, R. Masling, and B. N. Asmar, "Crowd dynamics discrete element multi-circle model," *Safety Science*, vol. 44, no. 5, pp. 395–417, Jun. 2006.
- [64] P. A. Thompson and E. W. March, "A computer model for the evacuation of large building populations," *Fire Safety Journal*, vol. 24, pp. 131–148, 1995.
- [65] M. Isobe, D. Helbing, and T. Nagatani, "Experiment, theory, and simulation of the evacuation of a room without visibility," *Phys. Rev. E*, vol. 69, p. 066132, Jun 2004. [Online]. Available: <http://link.aps.org/doi/10.1103/PhysRevE.69.066132>
- [66] R. Nagai, T. Nagatani, M. Isobe, and T. Adachi, "Effect of Exit Configuration on Evacuation of a Room Without Visibility," *Physica A: Statistical Mechanics and its Applications*, vol. 343, pp. 712–724, Nov. 2004.
- [67] J. M. Watts Jr, "Computer models for evacuation analysis," *Fire Safety Journal*, vol. 12, no. 3, pp. 237–245, 1987.
- [68] E. Kuligowski, "Review of 28 egress models," *Kuligowski (Eds.), Proceedings of the Workshop on Building Occupant Movement during Fire Emergencies*, pp. 68–90, 2004.
- [69] X. Zheng, T. Zhong, and M. Liu, "Modelling crowd evacuation of a building based on seven methodological approaches," *Building and Environment*, vol. 44, no. 3, pp. 437–445, Mar. 2009.

- [70] T. M. Kisko and R. L. Francis, “Evacnet+: a computer program to determine optimal building evacuation plans,” *Fire Safety Journal*, vol. 9, no. 2, pp. 211–220, 1985.
- [71] G. Lämmel, M. Rieser, and K. Nagel, “Large Scale Microscopic Evacuation Simulation,” *Pedestrian and Evacuation Dynamics 2008*, no. Chapter 48, pp. 547–553, Dec. 2009.
- [72] M. Bierlaire, G. Antonini, and M. Weber, “Behavioral dynamics for pedestrians,” in *International Conference on Travel Behavior Research*. Lucerne: EPFL, Aug. 2003.
- [73] P. Lino, G. Maione, and B. Maione, “Modelling and Simulation of Crowd Egress Dynamics in a Discrete Event Environment,” in *2009 IEEE Multi-conference on Systems and Control*. Saint Petersburg, Russia: IEEE, Jul. 2009, pp. 843–848.
- [74] L. Henderson, “On the fluid mechanics of human crowd motion,” *Transportation research*, vol. 8, pp. 509–515, 1974.
- [75] B. Banerjee, A. Abukmail, and L. Kraemer, “Advancing the Layered Approach to Agent-based Crowd Simulation,” in *22nd Workshop on Principles of Advanced and Distributed Simulation*. School of Computing, The University of Southern Mississippi, 2008, pp. 185–192.
- [76] ———, “Layered Intelligence for Agent-based Crowd Simulation,” *SIMULATION*, vol. 85, no. 10, pp. 621–633, Sep. 2009.
- [77] A. G. Hoekstra, J. Kroc, and P. M. A. Sloot, Eds., *Simulating Complex Systems by Cellular Automata*, ser. Understanding Complex Systems. Springer, 2010, iISBN: 978-3-642-12202-6.
- [78] W. Yuan and K. Tan, “An Evacuation Model Using Cellular Automata,” *Physica A: Statistical Mechanics and its Applications*, vol. 384, no. 2, pp. 549–566, Oct. 2007.
- [79] K. Yamamoto, S. Kokubo, and K. Nishinari, “Simulation for Pedestrian Dynamics by Real-Coded Cellular Automata (RCA),” *Physica A: Statistical Mechanics and its Applications*, vol. 379, no. 2, pp. 654–660, Jun. 2007.
- [80] W. Klein, G. Köster, and A. Meister, “Towards the calibration of pedestrian stream models,” in *Proceedings of the 8th international conference on Parallel processing and applied mathematics: Part II*, ser. PPAM’09. Berlin, Heidelberg: Springer-Verlag, 2010, pp. 521–528. [Online]. Available: <http://dl.acm.org/citation.cfm?id=1893586.1893650>

- [81] S. Marconi, “Mesoscopical modelling of complex systems,” Ph.D. dissertation, Université de Genève, Genève, 2002.
- [82] S. Marconi and B. Chopard, “A multiparticle lattice gas automata model for a crowd,” in *Cellular Automata*, ser. Lecture Notes in Computer Science, S. Bandini, B. Chopard, and M. Tomassini, Eds. Springer Berlin Heidelberg, 2002, vol. 2493, pp. 231–238. [Online]. Available: http://dx.doi.org/10.1007/3-540-45830-1_22
- [83] K. Nishinari, A. Kirchner, A. Namazi, and A. Schadschneider, “Extended floor field ca model for evacuation dynamics,” *IEICE Transactions on information and systems*, vol. 87, no. 3, pp. 726–732, 2004.
- [84] C. Henein and T. White, “Information in crowds: The swarm information model,” in *Cellular Automata*, ser. Lecture Notes in Computer Science, S. El Yacoubi, B. Chopard, and S. Bandini, Eds. Springer Berlin Heidelberg, 2006, vol. 4173, pp. 703–706. [Online]. Available: http://dx.doi.org/10.1007/11861201_83
- [85] Q. Zhang and B. Han, “Simulation Model of Pedestrian Interactive Behavior,” *Physica A: Statistical Mechanics and its Applications*, vol. 390, no. 4, pp. 636–646, Feb. 2011.
- [86] S. Bandini, M. L. Federici, and G. Vizzari, “Situated Cellular Agents Approach to Crowd modelling and Simulation,” *Cybernetics and Systems*, vol. 38, no. 7, pp. 729–753, Aug. 2007.
- [87] J. M. Epstein, “Agent-based computational models and generative social science,” *Complexity*, vol. 4, no. 5, pp. 41–60, May 1999.
- [88] E. Bonabeau, “Agent-based modelling methods and techniques for simulating human systems,” in *Arthur M. Sackier Colloquium of the National Academy of Sciences*. Icosystem Corporation, 545 Concord Avenue, Cambridge, May 2002, pp. 7280–7287.
- [89] X. Li, W. Mao, D. Zeng, and F.-Y. Wang, “Agent-based social simulation and modelling in social computing,” in *Intelligence and Security Informatics 2008 Workshops*. Chinese Academy of Sciences, 2008, pp. 401–412.
- [90] M. Wooldridge and N. R. Jennings, “Intelligent agents: theory and practice,” *The Knowledge Engineering Review*, vol. 10, pp. 115–152, 6 1995. [Online]. Available: http://journals.cambridge.org/article_S0269888900008122
- [91] M. Georgeff, B. Pell, M. Pollack, M. Tambe, and M. Wooldridge, “The belief-desire-intention model of agency,” in *Intelligent Agents V: Agents Theories*,

- Architectures, and Languages*, ser. Lecture Notes in Computer Science, J. MÃijller, A. Rao, and M. Singh, Eds. Springer Berlin Heidelberg, 1999, vol. 1555, pp. 1–10. [Online]. Available: http://dx.doi.org/10.1007/3-540-49057-4_1
- [92] E.-W. Augustijn-Beckers, J. Flacke, and B. Retsios, “Investigating the Effect of Different Pre-Evacuation Behavior and Exit Choice Strategies Using Agent-Based modelling,” *Procedia Engineering*, vol. 3, pp. 23–35, 2010.
- [93] A. H. Maslow, “A theory of human motivation.” *Psychological review*, vol. 50, pp. 370–396, 1943.
- [94] A. Orton, G. L. Clore, and A. Collins, *The Cognitive Structure of Emotions*. University of Illinois at Urbana-Champaign: Cambridge University Press, 1990.
- [95] M. Owen, E. R. Galea, and P. J. Lawrence, “The Exodus Evacuation Model Applied To Building Evacuation Scenarios,” *Journal of Fire Protection Engineering*, vol. 8, no. 2, pp. 65–84, Jan. 1996.
- [96] S. M. V. Gwynne, E. R. Galea, P. J. Lawrence, and L. Filippidas, “Modelling occupant interaction with fire conditions using the buildingEXODUS evacuation model,” *Fire Safety Journal*, vol. 36, pp. 327–357, 2001.
- [97] C. Hollmann, P. J. Lawrence, and E. R. Galea, “Introducing Emotion Modelling to Agent-Based Pedestrian Circulation Simulation,” in *Pedestrian and Evacuation Dynamics 2010*. Maryland, USA: University of Greenwich, Mar. 2010, pp. 1–4.
- [98] X. Pan, C. S. Han, K. Dauber, and K. H. Law, “A Multi-Agent Based Framework for the Simulation of Human and Social Behaviors During Emergency Evacuations,” *AI & SOCIETY*, vol. 22, no. 2, pp. 113–132, Jun. 2007.
- [99] S. Russel and P. Norvig, “Intelligent Agents,” in *Artificial Intelligence: A Modern Approach*. Prentice-Hall, Inc., 1995, pp. 30–52.
- [100] E. C. Tolman, “Cognitive maps in rats and men,” *The Psychological Review*, vol. 55, no. 4, pp. 189–208, 1948.
- [101] N. Hu, M. Lees, S. Zhou, and V. Viswanathan, “Pattern Based Motion For Crowd Simulation,” *Transactions on Edutainment*, pp. 1–10, 2011.
- [102] S. J. Guy, J. Chhugani, C. Kim, N. Satish, M. C. Lin, D. Manocha, and P. Dubey, “ClearPath: highly parallel collision avoidance for multi-agent simulation,” in *SCA '09: Proceedings of the 2009 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*. ACM Request Permissions, Aug. 2009.

- [103] S. J. Guy, J. van den Berg, and M. C. Lin, “Geometric methods for multi-agent collision avoidance,” in *Symposium on Computer Graphics*, University of North Carolina. Utah, USA: Proceedings of the 2010 annual symposium on Computational geometry, Jun. 2010, pp. 115–116.
- [104] H. Yeh, S. Curtis, S. Patil, and J. van den Berg, “Composite agents,” in *Symposium on Computer Animation*. University of North Carolina, 2008.
- [105] R. d. S. França, M. d. G. B. Marietto, and M. B. Steinberger, “Proposing a Cognitive Multi-Agent Model for the Panic in Crowds Phenomenon,” in *Applications of Digital Information and Web Technologies, 2009. ICADIWT '09. Second International Conference on the*, 2009, pp. 737–742.
- [106] V. Viswanathan and M. Lees, “An Information-Based Perception Model for Agent-Based Crowd and Egress Simulation,” in *2011 International Conference on Cyberworlds (CW)*, Jul. 2011, pp. 38–45.
- [107] —, “An information processing based model of pre-evacuation behavior for agent based egress simulation,” in *Pedestrian and Evacuation Dynamics 2012*, U. Weidmann, U. Kirsch, and M. Schreckenberg, Eds. Springer International Publishing, 2014, pp. 125–133. [Online]. Available: http://dx.doi.org/10.1007/978-3-319-02447-9_8
- [108] V. Viswanathan, C. Lee, M. Lees, S. Cheong, and P. Sloot, “Quantitative comparison between crowd models for evacuation planning and evaluation,” *The European Physical Journal B*, vol. 87, no. 2, 2014. [Online]. Available: <http://dx.doi.org/10.1140/epjb/e2014-40699-x>
- [109] J. Hochberg and E. McAlister, “A Quantitative Approach to Figural "Goodness",” *Journal of Experimental Psychology*, vol. 46, no. 5, pp. 361–364, Nov. 1953.
- [110] N. Cowan, “The Magical Number 4 in Short-Term Memory: a Reconsideration of Mental Storage Capacity,” *Behavioral and Brain Sciences*, vol. 24, no. 01, pp. 87–114, Feb. 2001.
- [111] L. Itti and C. Koch, “Computational modelling of visual attention,” *Nature*, vol. 2, pp. 194–203, Mar. 2001.
- [112] J. K. O'Regan, R. A. Rensink, and J. J. Clark, “Change-blindness as a result of 'mudsplashes',” *Nature*, vol. 398, p. 34, Mar. 1999.
- [113] J. Triesch, D. H. Ballard, M. M. Hayhoe, and B. T. Sullivan, “What you see is what you need,” *Journal of Vision*, vol. 3, pp. 1–9, Feb. 2003.

- [114] D. E. Broadbent, “Applications of Information Theory and Decision Theory to Human Perception and Reaction,” *Progress in Brain Research*, vol. 17, pp. 309–320, 1965.
- [115] R. W. Hill Jr, “Modelling Perceptual Attention in Virtual Humans,” in *Computer Generated Forces and Behavioral Representation*, Orlando, May 1999, pp. 1–11.
- [116] N. Courty, E. Marchand, and B. Arnaldi, “A new application for saliency maps: Synthetic vision of autonomous actors,” in *Image Processing, 2003. ICIP 2003. Proceedings. 2003 International Conference on*, vol. 3. IEEE, 2003, pp. 1065–1068.
- [117] Y. Kim, M. van Velsen, and R. W. Hill Jr, “Modeling dynamic perceptual attention in complex virtual environments,” in *Intelligent Virtual Agents*, ser. Lecture Notes in Computer Science, T. Panayiotopoulos, J. Gratch, R. Aylett, D. Ballin, P. Olivier, and T. Rist, Eds., vol. 3661. Springer Berlin Heidelberg, 2005, pp. 266–277. [Online]. Available: http://dx.doi.org/10.1007/11550617_23
- [118] H. Grillon and D. Thalmann, “Simulating gaze attention behaviors for crowds,” *Computer Animation and Virtual Worlds*, vol. 20, no. 2-3, pp. 111–119, Jun. 2009.
- [119] S. Okazaki and S. Matsushita, “A study of simulation model for pedestrian movement with evacuation and queuing,” *Engineering for Crowd Safety*, 1993.
- [120] J. Ondřej, J. Pettré, A.-H. Olivier, and S. Donikian, “A Synthetic-Vision Based Steering Approach for Crowd Simulation,” *ACM Transactions on Graphics (TOG) - Proceedings of ACM SIGGRAPH 2010*, vol. 29, no. 4, Jul. 2010.
- [121] J. van den Berg, S. J. Guy, and M. C. Lin, “Reciprocal n-Body Collision Avoidance,” *Robotics Research*, 2011.
- [122] J. van den Berg, M. C. Lin, and D. Manocha, “Reciprocal Velocity Obstacles for real-time multi-agent navigation,” in *Robotics and Automation, 2008. ICRA 2008. IEEE International Conference on*, 2008, pp. 1928–1935.
- [123] A. Kamphuis and M. H. Overmars, “Finding paths for coherent groups using clearance,” in *Proceedings of the 2004 ACM SIGGRAPH Symposium on Computer Animation*. Utrecht University, 2004.
- [124] J. van den Berg, S. Patil, J. Sewall, D. Manocha, and M. C. Lin, “Interactive Navigation of Multiple Agents in Crowded Environments,” in *2008 symposium on Interactive 3D graphics and games*, 2008, pp. 139–147.

- [125] Q. Song and N. Kasabov, "Ecm - a novel on-line, evolving clustering method and its applications," in *In M. I. Posner (Ed.), Foundations of cognitive science*. The MIT Press, 2001, pp. 631–682.
- [126] D. Wilkie, J. van den Berg, and D. Manocha, "Generalized velocity obstacles," in *Intelligent Robots and Systems, 2009. IROS 2009. IEEE/RSJ International Conference on*, 2009, pp. 5573–5578.
- [127] M. Whittle, *Gait Analysis: An Introduction*, 4th ed., M. Whittle, Ed. Butterworth Heinemann, Dec. 2006.
- [128] S. Nowak and A. Schadschneider, "Quantitative analysis of pedestrian counterflow in a cellular automaton model," *Phys. Rev. E*, vol. 85, p. 066128, Jun 2012. [Online]. Available: <http://link.aps.org/doi/10.1103/PhysRevE.85.066128>
- [129] N. Hu, M. H. Lees, and S. Zhou, "A pattern-based modelling framework for simulating human-like pedestrian steering behaviors," in *Proceedings of the 19th ACM Symposium on Virtual Reality Software and Technology*, ser. VRST '13. New York, NY, USA: ACM, 2013, pp. 179–188. [Online]. Available: <http://doi.acm.org/10.1145/2503713.2503723>
- [130] W. Daamen and S. P. Hoogendoorn, "Controlled Experiments to Derive Walking Behaviour," *Journal of Transport and Infrastructure Research*, vol. 3, no. 1, pp. 1–19, 2003.
- [131] S. D. Moeser, "Cognitive mapping in a complex building," *Environment and Behavior*, vol. 20, no. 1, pp. 21–49, 1988. [Online]. Available: <http://eab.sagepub.com/content/20/1/21.abstract>
- [132] K. Lynch, *The image of the city*. The M.I.T. Press, Cambridge, MA, 1960.
- [133] P. W. Thorndyke and B. Hayes-Roth, "Differences in spatial knowledge acquired from maps and navigation," *Cognitive Psychology*, vol. 14, no. 4, pp. 560 – 589, 1982. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/0010028582900196>
- [134] B. Kuipers, "Modelling spatial knowledge." *Cognitive Science*, vol. 2, no. 2, pp. 129–153, 1978.
- [135] A. W. Siegel and S. H. White, "The development of spatial representations of large-scale environments." *Advances in child development and behavior*, vol. 10, pp. 9–55, 1975.
- [136] B. Kuipers, D. G. Tecuci, and B. J. Stankiewicz, "The skeleton in the cognitive map: A computational and empirical exploration," *Environment*

- and Behavior*, vol. 35, no. 1, pp. 81–106, 2003. [Online]. Available: <http://eab.sagepub.com/content/35/1/81.abstract>
- [137] S. Gopal, R. L. Klatzky, and T. R. Smith, “Navigator: A psychologically based model of environmental learning through navigation,” *Journal of Environmental Psychology*, vol. 9, no. 4, pp. 309 – 331, 1989. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S027249448980012X>
- [138] C. Hölscher, S. J. BHüchner, T. Meilinger, and G. Strube, “Map use and wayfinding strategies in a multi-building ensemble.” in *Spatial Cognition*, ser. Lecture Notes in Computer Science, vol. 4387. Springer, 2006, pp. 365–380.
- [139] B. J. Stankiewicz, G. E. Legge, J. S. Mansfield, and E. J. Schlicht, “Lost in virtual space: Studies in human and ideal spatial navigation.” *Journal of Experimental Psychology: Human Perception and Performance*, vol. 32, no. 3, p. 688, 2006.
- [140] B. J. Stankiewicz and A. A. Kalia, “Acquisition of structural versus object landmark knowledge.” *Journal of Experimental Psychology: Human Perception and Performance*, vol. 33, no. 2, p. 378, 2007.
- [141] N. W. Bode and E. A. Codling, “Human exit route choice in virtual crowd evacuations,” *Animal Behaviour*, vol. 86, no. 2, pp. 347 – 358, 2013. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0003347213002443>
- [142] X. Pengfei, M. Lees, N. Hu, and V. Viswanathan, “Validation of Agent-Based Simulation Through Human Computation : An Example of Crowd Simulation,” in *12th International Workshop on Multi-Agent-Based Simulation*. Taipei, Taiwan: School of Computer Engineering, NTU, May 2011, pp. 1–12.
- [143] D. R. Michael and S. L. Chen, *Serious Games: Games That Educate, Train, and Inform*. Muska & Lipman/Premier-Trade, 2005.
- [144] D. Waller, J. M. Loomis, R. G. Golledge, and A. C. Beall, “Place learning in humans: The role of distance and direction information,” *Spatial Cognition and Computation*, vol. 2, pp. 333–354, May 2002.
- [145] D. R. Montello, M. Hegarty, A. E. Richardson, and D. Waller, “Spatial Memory of Real Environments, Virtual Environments, and Maps,” in *Human Spatial Memory: Remembering Where*, G. L. Allen, Ed. Lawrence Erlbaum Associates, Publishers, 2004, pp. 251–285.
- [146] N. Pelechano, C. Stocker, J. M. Allbeck, and N. I. Badler, “Being a Part of the Crowd: Towards Validating VR Crowds Using Presence,” in *Multiple values selected*, Estoril, Portugal, Jan. 2008, pp. 1–7.

- [147] N. W. F. Bode, A. U. Kemloh Wagoum, and E. A. Codling, "Human responses to multiple sources of directional information in virtual crowd evacuations," *Journal of The Royal Society Interface*, vol. 11, no. 91, 2014. [Online]. Available: <http://rsif.royalsocietypublishing.org/content/11/91/20130904.abstract>
- [148] R. T. Hays, J. W. Jacobs, C. Prince, and E. Salas, "Flight simulator training effectiveness: A meta-analysis," *Military Psychology*, vol. 4, no. 2, pp. 63–74, 1992.
- [149] S. Barry Issenberg, W. C. McGaghie, I. R. Hart, J. W. Mayer, J. M. Felner, E. R. Petrusa, R. A. Waugh, D. D. Brown, R. R. Safford, I. H. Gessner, D. L. Gordon, and G. A. Ewy, "Simulation technology for health care professional skills training and assessment," *The Journal of the American Medical Association*, vol. 282, no. 9, pp. 861–866, 1999. [Online]. Available: <http://dx.doi.org/10.1001/jama.282.9.861>
- [150] D. O. Hebb, *The Organization of Behavior*. John Wiley & Sons Inc, 1949.
- [151] A. D. Baddeley and G. J. Hitch, "Working memory," in *The psychology of learning and motivation*, G. Bower, Ed. New York: Academic Press, 1974, pp. 47–89.
- [152] E. Lindberg and T. Gärling, "Acquisition of locational information about reference points during locomotion with and without a concurrent task: Effects of number of reference points," *Scandinavian Journal of Psychology*, vol. 22, no. 1, pp. 109–115, 1981.
- [153] S. Garden, C. Cornoldi, and R. H. Logie, "Visuo-spatial working memory in navigation," *Applied Cognitive Psychology*, vol. 16, no. 1, pp. 35–50, 2002. [Online]. Available: <http://dx.doi.org/10.1002/acp.746>
- [154] T. Meilinger, M. Knauff, and H. H. Bülthoff, "Working memory in wayfinding: a dual task experiment in a virtual city," *Cognitive Science*, vol. 32, no. 4, pp. 755–770, 2008.
- [155] M. J. Sholl and S. K. Fraone, "Visuospatial Working Memory for Different Scales of Space: Weighing the Evidence," in *Human Spatial Memory: Remembering Where*, G. L. Allen, Ed. Lawrence Erlbaum Associates, Publishers, 2004, pp. 67–100.
- [156] T. Ishikawa and D. R. Montello, "Spatial knowledge acquisition from direct experience in the environment: Individual differences in the development of metric knowledge and the integration of separately learned places," *Cognitive Psychology*, vol. 52, no. 2, pp. 93 – 129, 2006. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0010028505000733>

- [157] G. Best, "Direction finding in large buildings," in *Architectural Psychology- Proceedings of the conference at Dalandhui, RIBA, London*, 1970, pp. 72–75.
- [158] J. Weisman, "Evaluating architectural legibility: Way-finding in the built environment," *Environment and Behavior*, vol. 13, no. 2, pp. 189–204, 1981. [Online]. Available: <http://eab.sagepub.com/content/13/2/189.abstract>
- [159] T. Gärling, E. Lindberg, and T. Mäntylä, "Orientation in buildings: Effects of familiarity, visual access, and orientation aids." *Journal of Applied Psychology*, vol. 68, no. 1, p. 177, 1983.
- [160] G. W. Evans and K. Pezdek, "Cognitive mapping: Knowledge of real-world distance and location information." *Journal of Experimental Psychology: Human Learning and Memory*, vol. 6, no. 1, p. 13, 1980.
- [161] M. J. O'Neill, "Effects of familiarity and plan complexity on wayfinding in simulated buildings," *Journal of Environmental Psychology*, vol. 12, no. 4, pp. 319 – 327, 1992. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0272494405800805>
- [162] B. J. Stankiewicz and K. Eastman, "Lost in virtual space ii: The role of proprioception and discrete actions when navigating with uncertainty," *Transactions on Applied Perception*, p. to appear, to appear.
- [163] J. McGonigal, *Reality is broken: Why games make us better and how they can change the world*. Penguin, 2011.
- [164] S. Kaplan and R. Kaplan, *Cognition and Environment: functioning of an uncertain world*. Ulrich's Bookstore, 1982.
- [165] Zyph, "Royal palace map," May 2011. [Online]. Available: <http://www.planetminecraft.com/project/royal-palace/>
- [166] M. Persson and J. Bergensten, "Minecraft," may 2009. [Online]. Available: <https://minecraft.net/>
- [167] SpaceManiac, "Permissionsbukkit," Sept 2012. [Online]. Available: <http://dev.bukkit.org/bukkit-plugins/permbukkit/>
- [168] V. Viswanathan, "Game analyser minecraft plugin," June 2013. [Online]. Available: <https://github.com/vaisaghvt/GameAnalyserMinecraftPlugin>
- [169] Z. Cheng, J. Caverlee, K. Lee, and D. Z. Sui, "Exploring Millions of Footprints in Location Sharing Services," in *Proceedings of the Fifth International Conference on Weblogs and Social Media*. Menlo Park, CA, USA: AAAI, Jul. 2011. [Online]. Available: <http://www.aaai.org/ocs/index.php/ICWSM/ICWSM11/paper/view/2783>

- [170] M. J. Spivey, M. Grosjean, and G. Knoblich, "Continuous attraction toward phonological competitors," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 102, no. 29, pp. 10 393–10 398, 2005. [Online]. Available: <http://www.pnas.org/content/102/29/10393.abstract>
- [171] J. B. Freeman and N. Ambady, "MouseTracker: software for studying real-time mental processing using a computer mouse-tracking method." *Behavior research methods*, vol. 42, no. 1, pp. 226–41, Feb. 2010. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/20160302>
- [172] S. Paris, J. Pettr, and S. Donikian, "Pedestrian reactive navigation for crowd simulation: a predictive approach," *Computer Graphics Forum*, vol. 26, no. 3, pp. 665–674, 2007. [Online]. Available: <http://dx.doi.org/10.1111/j.1467-8659.2007.01090.x>
- [173] A. Johansson, D. Helbing, H. Z. Al-abideen, and S. Al-bosta, "From crowd dynamics to crowd safety: A video-based analysis," *Advances in Complex Systems*, vol. 11, no. 04, pp. 497–527, 2008. [Online]. Available: <http://www.worldscientific.com/doi/abs/10.1142/S0219525908001854>
- [174] R. Hughes, "The flow of human crowds," *Annual Review Of Fluid Mechanics*, vol. 35, pp. 169–182, 2003.
- [175] Y. Tajima and T. Nagatani, "Scaling behavior of crowd flow outside a hall," *Physica A*, vol. 292, pp. 545–554, 2001.
- [176] M. Isobe, T. Adachi, and T. Nagatani, "Experiment and simulation of pedestrian counter flow," *Physica A: Statistical Mechanics and its Applications*, vol. 336, no. 3, pp. 638–650, 2004.
- [177] H. Xi, S. Lee, and Y.-J. Son, "An Integrated Pedestrian Behavior Model Based on Extended Decision Field Theory and Social Force Model," in *2010 Winter Simulation Conference*, Nov. 2010, pp. 824–836.
- [178] G. Peng and X. Ruihua, "3-Tier Architecture for Pedestrian Agent in Crowd Simulation," *Pedestrian and Evacuation Dynamics 2008*, no. Chapter 53, pp. 585–595, Dec. 2009.
- [179] D. Helbing, I. Farkas, and T. Vicsek, "Simulating dynamical features of escape panic," *Nature*, vol. 407, no. 6803, pp. 487–490, 2000.
- [180] P. Fiorini and Z. Shiller, "Motion planning in dynamic environments using the relative velocity paradigm," in *Robotics and Automation, 1993. Proceedings., 1993 IEEE International Conference on*, 1993, pp. 560–565.

- [181] S. J. Guy, M. C. Lin, and D. Manocha, “Modelling Collision Avoidance Behavior for Virtual Humans,” in *9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, Toronto, Canada, May 2010, pp. 575–582.
- [182] S. Luke, C. Cioffi-Revilla, L. Panait, K. Sullivan, and G. Balan, “MASON: A Multi-Agent Simulation Environment,” *Simulation : Transactions of the society for modelling and Simulation International*, vol. 82, no. 7, pp. 517–527, 2005.
- [183] J. Fruin, “Designing for pedestrians,” *Public Transportation United States*, 1992.
- [184] D. Bauer, “Comparing pedestrian movement simulation models for a crossing area based on real world data,” in *Pedestrian and Evacuation Dynamics*, R. D. Peacock, E. D. Kuligowski, and J. D. Averill, Eds. Springer US, 2011, pp. 547–556. [Online]. Available: http://dx.doi.org/10.1007/978-1-4419-9725-8_49
- [185] A. Seyfried and A. Schadschneider, “Fundamental diagram and validation of crowd models,” in *Cellular Automata*, ser. Lecture Notes in Computer Science, H. Umeo, S. Morishita, K. Nishinari, T. Komatsuzaki, and S. Bandini, Eds. Springer Berlin Heidelberg, 2008, vol. 5191, pp. 563–566. [Online]. Available: http://dx.doi.org/10.1007/978-3-540-79992-4_77
- [186] J. Lin, “Divergence measures based on the Shannon entropy,” *IEEE Transactions on Information Theory*, vol. 37, no. 1, pp. 145–151, 1991.
- [187] H. Abdi, D. Valentin, U. D. Bourgoigne, and B. Edelman, “Distatis: The analysis of multiple distance matrices,” in *Proceedings of the IEEE Computer Society: International Conference on Computer Vision and Pattern Recognition*, 2005, pp. 42–47.
- [188] X. Yang and Z. Wu, “Civilian monitoring video records for earthquake intensity: a potentially unbiased online information source of macro-seismology,” *Natural Hazards*, vol. 65, no. 3, pp. 1765–1781, 2013. [Online]. Available: <http://dx.doi.org/10.1007/s11069-012-0447-3>
- [189] X. Yang, Z. Wu, and Y. Li, “Difference between real-life escape panic and mimic exercises in simulated situation with implications to the statistical physics models of emergency evacuation: The 2008 wenchuan earthquake,” *Physica A: Statistical Mechanics and its Applications*, vol. 390, no. 12, pp. 2375 – 2380, 2011. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0378437110008770>
- [190] J. McKellar and A. Fettiq, *Twisted: Network Programming Essentials*, 2nd ed. O’Reilly, 2013.

- [191] A. Teo, V. Viswanathan, M. Lees, and C. Wentong, “Analysing the Effectiveness of Wearable Wireless Sensors in Controlling Crowd Disasters,” in *2014 International Conference on Computational Science (ICCS)*, Jul. 2014.
- [192] H. Aydt, S. J. Turner, W. Cai, and M. Y. H. Low, “Symbiotic Simulation Systems: An Extended Definition Motivated by Symbiosis in Biology,” in *2008 ACM/IEEE/SCS Workshop on Principles of Advanced and Distributed Simulation (PADS)*. IEEE, 2008, pp. 109–116.