12-12-2011

Pedestrian Leadership and Egress Assistance Simulation Environment (PLEASE)

Kyle D. Feuz
Utah State University

PEDESTRIAN LEADERSHIP AND EGRESS ASSISTANCE SIMULATION ENVIRONMENT (PLEASE)

by

Kyle D. Feuz

A thesis submitted in partial fulfillment
of the requirements for the degree

of

MASTER OF SCIENCE

in

Computer Science

Approved:

Dr. Vicki Allan Major Professor	Dr. Daniel Bryce Committee Member
Dr. Curtis Dyreson Committee Member	Dr. Mark R. McLellan Vice President for Research and Dean of the School of Graduate Studies

UTAH STATE UNIVERSITY
Logan, Utah

2011

ABSTRACT

Pedestrian Leadership and Egress Assistance Simulation Environment (PLEASE)

by

Kyle D. Feuz, Master of Science

Utah State University, 2011

Major Professor: Dr. Vicki Allan
Department: Computer Science

Over the past decade, researchers have been developing new ways to model pedestrian egress especially in emergency situations. The traditional methods of modeling pedestrian egress, including flow-based modeling and cellular automata, have been shown to be poor models of human behavior at an individual level, as well as failing to capture many important group social behaviors of pedestrians. This has led to the exploration of agent-based modeling for crowd simulations including those involving pedestrian egress. In this work, we implement a multi-agent simulation model to specifically address the issues of pedestrian route selection with uncertainty and group formation. Using this model, we evaluate different heuristic functions for predicting good egress routes for a variety of real building layouts. We also introduce reinforcement learning as a means to represent individualized pedestrian route knowledge. Finally, we implement a group formation technique, which allows pedestrians in a group to share route knowledge and reach a consensus in route selection. Using the group formation technique, we consider the effects such knowledge sharing and consensus mechanisms have on pedestrian egress times.

(83 pages)

PUBLIC ABSTRACT

Pedestrian simulation models are used in many different applications including, the design of safer buildings, the validation of fire codes, and automatic video surveillance and tracking. By improving the simulation model used, each of the application areas can experience similar improvements in accuracy. Current simulation models fail to address key concerns in representing pedestrian knowledge and in accurately modeling group formation. This project has at its core the goal of bringing attention these areas of concern and providing an initial look at ways to solve these problems.

The Pedestrian Leadership and Egress Assistance Simulation Environment (PLEASE) is developed specifically to allow for greater diversity in representing pedestrian knowledge without requiring a large time commitment from the end-user of the simulation. PLEASE additionally provides new mechanism for representing group formation, information sharing, and collaborative decision making. By providing these features, PLEASE can help bring about more realistic simulation models for the design of safer buildings, the development of better fire codes, and the implementation of more accurate video surveillance and tracking models.

To my loving wife Shara

ACKNOWLEDGMENTS

This thesis would not have been possible without the invaluable help of my major professor, Dr. Vicki Allan. I have known and worked with her for several years, beginning as an undergraduate researcher and now as a master's student. Throughout this time, she has always been a great support to me. Her advice, encouragement, constructive criticism, and support have pushed me to be the best I can be and have helped me get where I am today. She has trusted me with responsibilities I did not know I was ready for, but she knew I was. To be short, I owe a debt of gratitude to Dr. Vicki Allan and will always be grateful for the help she has given me.

Along with Dr. Allan, there are many others whose support has been essential to the completion of this thesis. My committee members and the CS department faculty and staff have also helped a great deal by providing me with the tools and techniques necessary to succeed. Rather than list specific individuals who have helped me along the way, I have grouped them together to avoid forgetting someone but to each of you I give my sincere thanks.

I would be truly ungrateful if I did not include, in these acknowledgments, all of my family, but especially my parents who have taught me the importance of hard work, honesty and other core values. Their guidance has allowed me to follow the path I have chosen, and their love and encouragement has helped me whenever I have felt inadequate.

Two other very important people need to be given credit here as well. They are my wife and daughter. These two have been by my side throughout the entire process. This thesis has probably been harder on them than it has been on me, but they have loved me without fail. I would not have completed this thesis were it not for the constant and unwavering support I feel from them. They patiently listened on countless occasions as I explained what I was working on despite their lack of interest in the topic. They have been my biggest supporters, and I am truly grateful for their love and support.

Finally, I must also acknowledge the help of my Heavenly Father who makes all things possible. At times, my faith in Him is all that kept me going, and He always provided the boost I needed to overcome whatever obstacles I faced.

This work was funded by an NSF grant, COAL: Coalition Organization with Agent Leadership NSF Proposal ID: 0812039.

Kyle D. Feuz

CONTENTS

LIST OF TABLES

Table Page

LIST OF FIGURES

CHAPTER 1

INTRODUCTION

Pedestrian egress has become an important research area over the last few decades. One of the chief concerns in this area is providing a means to develop safe and efficient egress routes for pedestrians in an emergency. To do this effectively, first one needs to understand human behavior and reactions; second, one must be able to design effective egress routes around these characteristics while also considering the effect of information content in the form of exits signs, maps, and egress assistance devices. Real-world experiments are too dangerous and too expensive to be a practical way of learning about egress efficiency. For this reason, simulation models have been developed to demonstrate crowd behavior in an emergency. In addition to the application of the simulation model discussed above, pedestrian simulation models have been applied to a wide range of applications. A few such applications include the design of buildings to detect potential bottlenecks and the determination and justification of fire-code regulations [24]. A novel application of pedestrian simulators that has recently appeared in research uses the simulator in conjunction with computer-vision techniques to produce a more accurate automated pedestrian tracking system [1]. In the past, stochastic models, flow-based models, cellular automata, or particle systems have been used to model pedestrian flow during an emergency evacuation [48, 57]. While each of these methods has its strengths, none of these models sufficiently captures the individual behavior of a pedestrian while accurately modeling the behavior of the crowd as a whole [27]. A new paradigm is needed to capture such complex actions. Thus, researchers turn to agent-based modeling.

In agent-based modeling, agents act autonomously to affect the system as a whole. The sophistication of the agent can vary from a reactive agent that acts according to a fixed set of rules to more advanced agents such as planning and learning agents. One of the biggest challenges to agent-based modeling for pedestrian egress involves deciding how the agent should make decisions. A model that is too simple will not produce an accurate simulation of human behavior. On the other hand, a model that is too complex will be difficult to use and lack scalability, thereby reducing the utility of the model.

In this paper, we develop a new multi-agent simulation model, Pedestrian Leadership and Egress Assistance Simulation Environment (PLEASE), which addresses four critical challenges in pedestrian egress: How can pedestrians successfully navigate buildings in which they have no prior experience and, therefore, no knowledge about the building layout? How can pedestrians be given different amounts of knowledge about the building without requiring the user to manually enter the information? What effect does knowledge of congestion levels have upon pedestrian egress? And, what effect does group formation have on pedestrian egress when group members are allowed to share and communicate route information?

The layout for the remainder of this thesis is as follows. Chapter 2 describes the PLEASE system in more detail. In chapter 3, we cover pedestrian route selection when pedestrians have no prior knowledge of route costs. In chapter 4, we present a novel use of reinforcement learning to represent pedestrian knowledge and apply it to learning congestion costs. Chapter 5 discusses pedestrian group formation and its impact on egress times. Finally, chapter 6 summarizes our findings and present possible directions for future work. The reader should be aware that chapters 3 - 5 are papers that have been submitted to different conferences. As such, each chapter contains some redundant introductory material.

CHAPTER 2
SIMULATION ENVIRONMENT

PLEASE is built on top of the Multi-Agent Simulator of Neighborhoods (MASON) framework developed at George Mason University. MASON is a discrete-event multi-agent simulation library core in Java with additional 2D and 3D visualization tools [34]. MASON provides us with an interface for setting the model parameters, visualizing the simulation and controlling the stepping of simulation events. It does not provide the implementation for the actual pedestrian agents. Figure 2.1 shows sample screen-shots of the user-interface. In Figure 2.1(a), we see MASON controls for automating the simulation. Figure 2.1(b) shows MASON's ability to allow the system user to view and modify individual agent parameters. Finally, Figure 2.1(c) shows MASON's ability to allow the system user to view and modify simulation parameters.

Through the MASON framework, we implement the pedestrian agents used in PLEASE, as well as the environment in which the pedestrians interact. Figure 2.2 shows an example screen-shot of the pedestrian agents and building environment. The pedestrians are colored according to the group to which they belong. Building walls are shown in blue, and building signage is shown in green and gold.

All the pedestrian simulation agents in PLEASE use a two-tier navigational model. The first tier performs the high level function of route selection and path planning while the second tier performs the lower-level tactical navigation and collision avoidance. In this chapter, we first describe the implementation of the lower-level navigation model which uses social forces. We then describe the higher level route planning model and its various components. Finally, we introduce the dynamic coalition formation model pedestrian agents used to form cooperative groups. Splitting the pedestrian navigation model into two tiers is not a novel idea, nor is the formation of groups. However, the details of our high-level

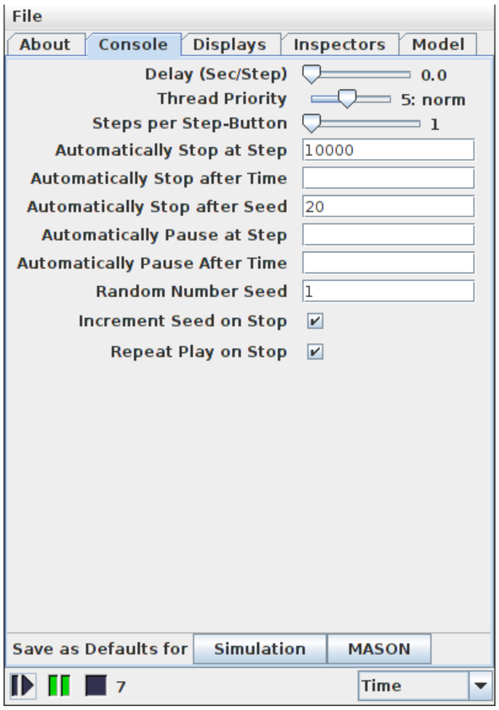

(a) Controls to automate the simulation

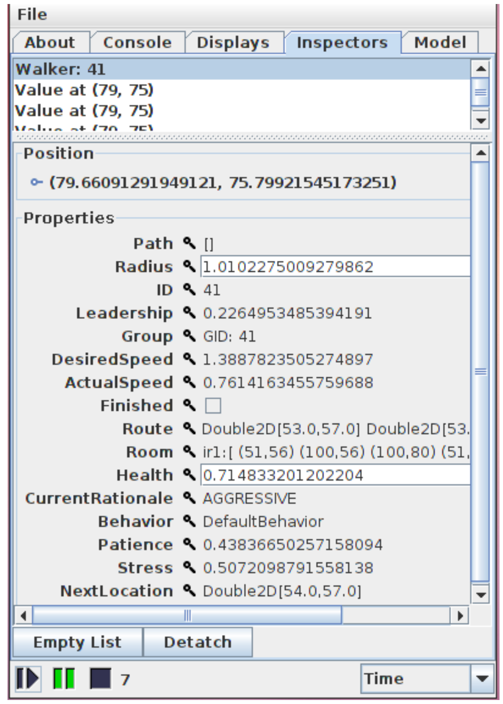

(b) Controls to view and manipulate agent parameters

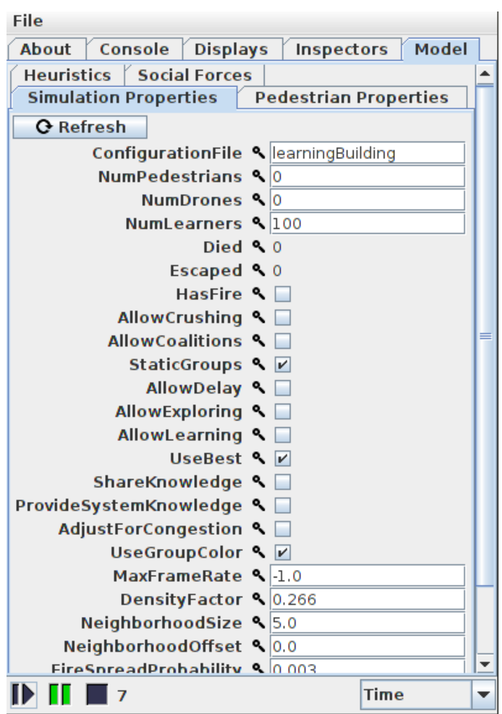

(c) Controls to view and manipulate simulation parameters

Figure 2.1: Example screen-shots of the user controls in PLEASE available through the MASON framework.

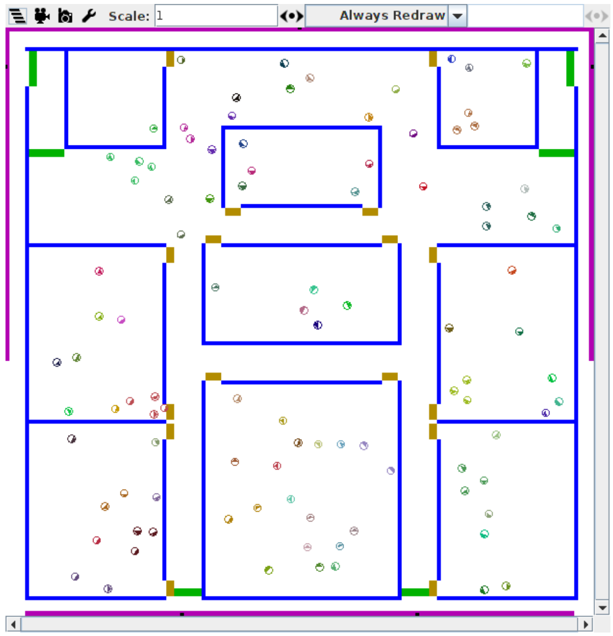

Figure 2.2: Screen-shot of the PLEASE simulator in action. Pedestrians in the same group have the same color. Walls are shown in blue and building signage is shown in green and gold.

navigation model and group formation techniques are a novel contribution to the field.

2.1 Low-Level Navigation

PLEASE uses the social force model to perform low-level tactical navigation and collision avoidance. The social force model is selected for the low-level tactical navigation for four main reasons: it is simple to understand and implement, it is widely used in many simulation models, it successfully reproduces many crowd phenomena, and it has been validated using actual pedestrian data [19, 29, 33, 35]. The social force model states that pedestrian movement can be approximated by applying multiple force vectors to a pedestrian. The force vectors can include attraction vectors for groups or goals and repulsion vectors for obstacles, fire, or other pedestrians. The concept first developed by Lewin [30] was later put into these mathematical terms by Helbing [17, 18] and has since been further enhanced in [19, 35]. For clarity, we define the forces which are applied to pedestrians here, but for a more complete discussion, we refer the reader to [19, 35].

The social force vector $\vec{f_i}(t)$ applied to the pedestrian i is the sum of the attraction forces $\vec{f_i}^{att}(t)$ and the repulsion forces $\vec{f_i}^{att}(t)$ plus a random component $\vec{f_i}^{rand}(t)$, as shown in Formula 2.1. Figure 2.3 depicts some example forces that may be applied to a pedestrian.

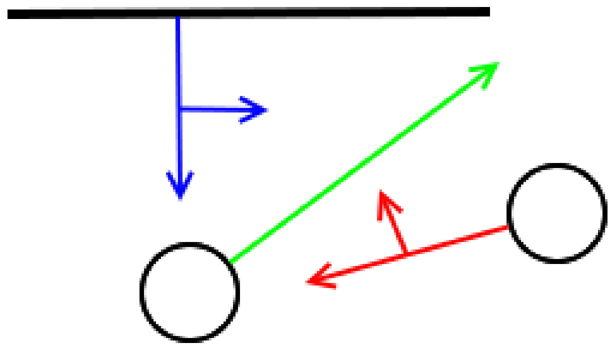

Figure 2.3: Graphical depiction of social force acting on a pedestrian. The green vector represents the driving force. The blue vector represent an obstacle repulsion force and the red vector represents a pedestrian repulsion force.

$$\vec{f_i}(t) = \vec{f_i}^{att}(t) + \vec{f_i}^{rep}(t) + \vec{f_i}^{rand}(t) \tag{2.1}$$

The attraction forces can be further broken down into a driving force, $\vec{f_i}^{drv}(t)$ (see Formula 2.3), in the direction of the goal and an attraction force, $\vec{f_i}^{grp}(t)$ (see Formula 2.4), towards the group center, as shown in Formula 2.2.

$$\vec{f_i}^{att}(t) = \vec{f_i}^{drv}(t) + \vec{f_i}^{grp}(t) \tag{2.2}$$

The driving force represents a pedestrian's desire to reach a destination while traveling at a certain speed. It is given by Formula 2.3, where v_i^0 is the desired speed of the pedestrian, $\vec{e}_i^0$ is the desired direction of the pedestrian, $\vec{v}_i$ is the actual velocity of the pedestrian, and τ_i is the relaxation time to adjust to the desired velocity.

$$\vec{f_i}^{drv}(t) = \frac{1}{\tau_i}(v_i^0\vec{e}_i^0 - \vec{v}_i) \tag{2.3}$$

The group force represents a pedestrian's desire to maintain close proximity to other members of its group. It is given by Formula 2.4, where β represents the strength of the interaction and $\vec{U_i}$ is the unit vector pointing from the agent to the group's center of mass. q has a value of 1 if the distance between pedestrian i and the group's center of mass is greater than $\frac{N-1}{2}$ meters, or a value of 0 otherwise where N is the size of the group. The threshold value $\frac{N-1}{2}$ is determined empirically in [35].

$$\vec{f}_i^{grp}(t) = q\beta\vec{U}_i \tag{2.4}$$

Similar to the attraction forces, the repulsion forces can be further broken down into repulsion forces from other pedestrians, $\vec{f}_{ij}(t)$ (see Formula 2.6), and repulsion forces from obstacles, $\vec{f}_i^{obs}(t)$ (see Formula 2.7), as shown in Formula 2.5.

$$\vec{f}_i^{rep}(t) = \vec{f}_{ij}(t) + \vec{f}_i^{obs}(t) \tag{2.5}$$

The repulsion force from other pedestrians is composed of three repulsive forces. An isotropic force $\vec{f}_{ij}^{iso}(t)$, an anisotropic force $\vec{f}_{ij}^{ani}(t)$, and a contact force $\vec{f}_{ij}^{con}(t)$. The isotropic force (directionally independent) is of equal value independent of the direction of the force and has a short range of effectiveness. The anisotropic force (directionally dependent) has greater weight when the pedestrians are in front of pedestrian, i, and has a longer range of effectiveness. The contact force is isotropic, but the range is limited to pedestrians close enough to be in physical contact. The exact specification of these forces are found in [19].

$$\vec{f}_{ij}(t) = \vec{f}_{ij}^{iso}(t) + \vec{f}_{ij}^{ani}(t) + \vec{f}_{ij}^{con}(t) \tag{2.6}$$

Similarly, the repulsion force from obstacles is composed of two forces, an avoidance force $\vec{f}_i^{avd}(t)$ and a contact force $\vec{f}_i^{con}(t)$.

$$\vec{f}_i^{obs}(t) = \vec{f}_i^{avd}(t) + \vec{f}_i^{con}(t) \tag{2.7}$$

2.2 High-Level Navigation: Route Selection and Path Planning

High-level pedestrian navigation involves both route selection and path planning. In order to understand the process an agent uses in route selection and path planning, we must first understand the representation for buildings that PLEASE uses and both the type and amount of knowledge the agent has.

PLEASE allows users to load building layouts that are stored in a special XML file. The XML file specifies where the rooms, corridors, and walls are located, as well as indicating the distribution of pedestrians in the different rooms. The XML file also identifies the decision points in the building. A *decision point* is defined as point in the building at which an agent must decide upon the next location in the route. These points may be placed at arbitrary locations, but typically decision points are placed at doorways and intersections. By only placing decision points at doorways and intersections, the placement of the decision point does not require the pedestrian to pass through an area which they would not normally pass through when navigating from one area of the building to another.

PLEASE provides agents with three types of knowledge: heuristic knowledge, learned knowledge, and system knowledge. A brief overview of each of these types of knowledge is given here. They are discussed in more detail in chapters 3 and 4.

Pedestrians using local knowledge assign heuristic costs to known or visible decision points, and then choose the decision point with the least cost. These costs are based upon a variety of weighted factors including distance, congestion levels, number of neighbors moving towards that location, corridor width, angles, and experience. The heuristics used and the weights applied to each heuristic can be specified by the user. This technique does not require any knowledge about the building to be given to the pedestrian agent other than knowledge that can be directly observed during the simulation.

Pedestrian agents can use previously learned costs to estimate the exit cost from a decision point. PLEASE allows the agents to be trained using reinforcement learning in a particular building before running the actual training. During the training, the agents may learn distance costs, congestion levels, or other useful heuristic factors. The amount of training an agent receives, and the type of information the agent learns can be specified by the user. Using reinforcement learning to simulate pedestrians with prior knowledge of a building is a novel application of the technique and does not require the user to manually enter individualized pedestrian knowledge.

Finally, PLEASE can also provide agents with complete route distance information and the current congestion levels along the route. This situation models a pedestrian with access to an egress assistance device capable of providing such knowledge and can be used to evaluate the effectiveness of such devices. Many simulation models provide agents with complete route distance information.. The inclusion of complete, dynamic congestion information is much less common, and to our knowledge, the effect such information has on pedestrian egress has not been previously considered.

2.3 Coalition Formation

In addition to the two-tier navigation model, we implement a group formation model, that allows pedestrians to form static and dynamic groups. Again, we provide a brief overview of the model here, but details are given in chapter 5. Static groups are only formed at the beginning of the simulation. As the name implies, once the static group has formed, the membership of the group does not change throughout the simulation. Static groups represent family or social bonds that exist between pedestrians. Dynamic groups, on the other hand, constantly change throughout the simulation. In PLEASE, dynamic groups form based upon the principles of satisficing utility theory [13,50,54]. If an agent's expected utility is below a certain threshold, the agent will seek to join or form a group. PLEASE is extensible, so the function used to calculate an agent's utility can be easily changed. In chapter 5, we describe utility functions based upon the stress level or knowledge level of an agent.

CHAPTER 3

PEDESTRIAN ROUTE SELECTION WITH IMPERFECT KNOWLEDGE[1]

3.1 Introduction

In everyday life, pedestrians frequently utilize heuristics to make potentially complicated decisions within an allotted time frame. These heuristics often allow a person to make a reasonable decision without requiring extensive amounts of time and energy. However, in some situations the heuristic or the data a person relies on may not be valid and can lead to poor decisions. In life threatening situations, the effective use of heuristics becomes even more valuable. One example arises from the evacuation of a burning building. Selecting the best route can mean the difference between life and death, but all too frequently, insufficient information is used in making such a decision. Currently, detailed information about pedestrian egress from actual emergencies is not widely available, and thus it is too costly, too dangerous, and impractical to determine what heuristics pedestrians use to select an egress route when in emergency situations. Pedestrian simulation models have been designed to address this problem.

Current pedestrian simulation models generally fall into one of two categories when determining pedestrian route selection. Either the model assumes pedestrians have perfect knowledge of the building layout and are thus able to select the best route, or the model assumes that pedestrians only know about the immediately visible routes and must make a decision based upon some heuristic [15, 39, 51]. Models that assume perfect knowledge are clearly unrealistic for many situations as not all pedestrians are familiar with any given building layout. Yet models relying upon a heuristic may also be unrealistic in the amount

[1] Feuz, Kyle and Allan, Vicki "Pedestrian Route Selection with Imperfect Knowledge" Condensed version accepted for publication in the 4^{th} International Conference on Agents and Artificial Intelligence 2012

of information provided to pedestrians. Typically such heuristic models consider distance, congestion, or social comparison for making the route selection decision. In this paper, we show that using distance as the sole means of comparing visible routes leads to poor egress times when the total route distances are unknown. In addition, it is actually equivalent to the shortest-leg first heuristic which Golledge found [12] to be less preferred than many other heuristics. Therefore, if total route distances are assumed to be unknown to pedestrians, other route selection heuristics must be used to supplant the unknown information.

Most models assume perfect distance information. According to Golledge, in observing heuristics used in practice by actual pedestrians, shortest distance is the most frequently used information. In traffic systems, perfect distance information is reasonable because road signs often indicate the distance to a desired destination. In egress systems, perfect distance information is not reasonable because there are no such distance signs and some exits may be obstructed. Pedestrians must rely upon the perceived distance of routes (which may not be accurate and cannot be known for previously unexplored or occluded routes). Therefore, the lack of perfect distance information needs to be addressed with other heuristics.

In this paper, we consider the effect of various heuristic functions on pedestrian egress time for a variety of different building layouts. Determining a comprehensive list of heuristic functions that pedestrians might use is outside of the scope of this paper. However, several common factors including distance, signage, corridor width, congestion/usage, and common consensus are used to produce a variety of realistic heuristic functions.

3.2 Related Work

In [38], Ozel raises several pertinent questions regarding the issue of stress management and the decision process. He suggests that pedestrians utilize various coping mechanisms (heuristics) to make a decision in a time-pressured environment. In particular, the familiarity of routes and the negative connotations of emergency exits are shown to have a large impact on the route choice of pedestrians in an emergency.

Hoogendoorn and Bovy use distance and congestion heuristics for their route-choice and activity scheduling model [23]. Gwynne et al. indicate that pedestrians maintain social roles and norms even during emergency situations [14]. This may lead to pedestrians choosing a different route based on the choices of their peers. Similarly, Fridman and Kaminka develop a pedestrian simulation model based upon social comparison theory [11].

The simulation models MASSEgress, Simulex, and buildingEXODUS each use distance as the primary factor in selecting an available exit [15, 25, 39, 51]. MASSEgress can be configured to use other factors when selecting a route, and the agents only select visible routes. Simulex assumes the agents know all routes and the total distance of each route. buildingEXODUS assumes the agents know about a set of user specified routes or all routes in the absence of the specification. In each simulation model, the total distance of each route considered is known by the pedestrian.

Golledge ranks the preference and prevalence of several different heuristic values pedestrians use in navigation and route selection in an outdoor environment, such as a college campus, through the use of questionares and observations [12]. Table 3.1 summarizes the rankings found by Golledge to be the most commonly used heuristics in pedestrian route selection. These rankings indicate that pedestrians prefer routes that are direct, quick, and easy to navigate, with some preference being given to routes that are more aesthetically pleasing. In emergency situations, one would expect the characteristics of directness and quickness to remain prominent while the importance of scenery is irrelevant. Golledge also found that the route selection criteria used differed for various route layouts and that a combination of multiple criteria may give better results. Frankenstein et al. create virtual representations of various corridors to determine the effect of geometry on pedestrian route selection and show that some of the same heuristics used for vehicle route planning are also applicable to pedestrian route planning [10].

Other researchers focusing on the egress problem are concerned with building design. Both Helbing and Hassan use modern evolutionary algorithm techniques to modify and generate building layouts that are more conducive to safe and efficient pedestrian egress [16, 19]. Primarily, the algorithms are concerned with reducing congestion, increasing flow rates, and decreasing the likelihood of trampling and crushing.

Table 3.1: Ranking of Criteria Preference in Route Selection from [12]

Criteria (Heuristic)	Rank
Shortest Distance	1
Least Time	2
Fewest Turns	3
Most Scenic	4
First Noticed	5
Longest Leg First	6
Many Curves	7
Many Turns	8
Different From Previous	9
Shortest-leg First	10

3.3 Heuristics

Our research evaluates a variety of heuristics.

3.3.1 Distance

Distance is a common heuristic used by pedestrians when making a route selection and is the primary heuristic in many search based strategies [12, 23]. Unless motivated by other factors, pedestrians will choose the route with the shortest distance [12]. The distance heuristic is inapplicable in the absence of complete route distance knowledge because missing data causes the distance relation to be a partial order.

3.3.2 Shortest-leg First

The shortest-leg first heuristic is a greedy distance heuristic. Rather than utilizing the total route distance, shortest-leg first selects the route with the shortest distance to the next decision point. Pedestrians only require distances to route goals that are within their field of view.

3.3.3 Angle

Another common heuristic used by pedestrians when selecting a route is the least angle heuristic [7, 22]. This heuristic selects the path that is closest in terms of angularity to a direct line between their current position and the final goal [4]. For example, in Fig. 3.1 the agent would select Route B when using the least angle heuristic because it is the route whose angle differs the least from the angle to the goal. Unfortunately, the end goal location(s) is not always known. In such situations, the least-angle heuristic cannot be applied.

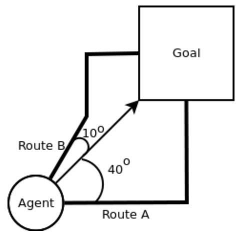

Figure 3.1: Depicts the values use to calculate the least-angle heuristic

3.3.4 Fewest Turns

We can modify the least angle heuristic to accommodate unknown goals by selecting the closest path in terms of angularity to the current walking orientation. This new heuristic is actually a greedy version of the fewest turns heuristic as it tries to minimize the number of direction changes an agent makes while exiting from the building.

3.3.5 Signage

Most buildings include navigational signs to assist pedestrians in locating their desired location [28]. These signs might include exit signs, emergency exits signs, room signs, or navigational maps. Pedestrians then use these signs to navigate effectively through the building. One might think that during an emergency evacuation, pedestrians would look for exit signs or emergency exit signs, but Ozel found that emergency exit signs often have negative connotation and are avoided even in emergency situations [38]. Exit signs, on the other hand, are commonly used by pedestrians in navigational planning when in an unfamiliar building. Room signs can also be used in egress route planning as they indicate that an egress route through a particular door is unlikely because individual rooms rarely contain egress routes.

3.3.6 Corridor Width

Main corridors tend to be preferred by pedestrians especially when navigating an unfamiliar building [21]. Like the relative width of various categories of roads, main corridors are wider than auxiliary hallways, so the width of the corridor or doorway can frequently be an indication of a main route of travel leading to an exit.

3.3.7 Congestion/Usage

The number of pedestrians currently gathered around a location may be either a positive or a negative indicator of a desirable exit route [14, 27, 39]. It may indicate the route is a good choice because others are using it. However, it may also indicate that another route should be considered to avoid the congestion. Intuition suggests that if pedestrians are unsure about the situation, they will follow others. If they are more confident and know multiple routes, they will seek alternatives to avoid congestion.

3.3.8 Common Consensus

Unless pedestrians are traveling in a group, they will not know the exact route of nearby pedestrians. However, by observing the velocity and past movement, an agent may predict the immediate destination of a neighboring pedestrian. Seeing agents who are traveling in the same direction may bolster the confidence level of an agent in choosing a particular route [38]. Traveling in the same direction as others is also easier as one does not have to fight against the general direction flow, improving the overall pedestrain flow.

3.3.9 Past Experience

Pedestrians will typically choose routes with which they are familiar especially when under time constraints [38]. Familiarity allows the pedestrian to feel more comfortable and confident in the route selected and allows the pedestrian to concentrate on other cognitive tasks. However, selecting a new route may or may not lead to a better solution. When the current known route cost is within an acceptable limit, pedestrians are unlikely to change.

The cost of a route may be based upon any number of factors such as distance, congestion, time, and so on.

3.4 Simulation Environment

This study is performed using the Pedestrian Leadership and Egress Assistance Simulation Environment (PLEASE). PLEASE is built upon the multi-agent modeling paradigm wherein each pedestrian is represented as an individually rational agent capable of perceiving the environment and reacting to it. In PLEASE, pedestrian agents can perceive obstacles, hazards, routes, and other agents. Agents use a two-tier navigational module to control their movement within the simulation environment. The high-level tier evaluates available routes and selects a destination goal. The low-level tier, based on the social force model [19], performs basic navigation and collision avoidance.

PLEASE implements several of the heuristics outlined in Section 3.3. Here, we describe the implementation details used for the heuristics of interest. Many simulation models assume that total route distance is known and/or the end goal location is known for at least some subset of the available routes. However, in this paper, we are interested in the case when no additional information, beyond what is immediately visible to the pedestrian, is provided. For this reason, the distance, angle, and past experience heuristics are not considered.

Each heuristic can potentially be combined with any other heuristic. To facilitate the integration of multiple heuristics, all heuristic values are normalized so that the unweighted cost falls between 0 and 1.

The leg cost of route x, $L(x)$, is given by Formula 3.1 where w_l is the weight applied to the shortest-leg heuristic, and $d_{i,x}$ is the distance from agent i to route goal x. The distance is normalized using the maximum distance between two points on the simulation map. Agents in the simulation are able to accurately estimate the distance to visible points within the simulation model. The distance to locations which are occluded by walls or other obstacles cannot be estimated without prior knowledge.

$$L(x) = \frac{w_l * d_{i,x}}{(maxDistance)} \tag{3.1}$$

The turn cost of route x, $T(x)$, is given by Formula 3.2, where w_t is the weight applied to the fewest turns heuristic, and $a_{i,x}$ is the angle in radians between the orientation of agent i and the direction to route goal x from agent i. π acts as the normalization factor since no angle will be greater than π. See Figure 3.2 for an example.

$$T(x) = w_t * a_{i,x}/\pi \tag{3.2}$$

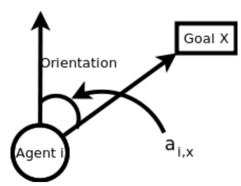

Figure 3.2: Depicts the values used to calculate the angle heuristic.

The signage cost of route x, $S(x)$, is given by Formula 3.3, where w_s is the weight applied to the signage heuristic, and $getExitWeight(x)$ is the cost associated with the given signage value. getExitWeight is a simple lookup table that takes the signage of a route and looks up the user-specified cost of that sign. A user may enter any arbitrary sign cost, but by default, PLEASE uses the cost shown in Table 3.2. To be consistent with the other heuristics, cost should be specified as a value between 0 and 1.

$$S(x) = w_s * getExitWeight(x) \tag{3.3}$$

Table 3.2: Signage Costs Used by Default in PLEASE.

Sign Type	Cost
Emergency Exit	0.25
Exit	0.0
Room	1.0
None	0.5

The simple signage cost of route x, $SS(x)$, is given by Formula 3.4, where w_{ss} is the weight applied to the simple signage heuristic, and $getSimpleSignage(x)$ is the cost associated with the simple signage of route x. A route has a simple signage cost of 0 if the doorway of the route is a direct exit and is marked with an exit sign. All other routes have a simple signage cost of 1. This causes a pedestrian to ignore all building signage except for exit signs over direct exits. The simple signage heuristic allows us to compare the effect that different amounts of building signage and different levels of attention to the building signage have upon the egress times of pedestrians.

$$SS(x) = w_{ss} * getSimpleSignage(x) \tag{3.4}$$

The width cost of route x, $W(x)$, is given by Formula 3.5, where w_w is the weight applied to the corridor width heuristic, and $getCorridorWeight(x)$ is the cost assigned to the given corridor width. The getCorridorWeight is a lookup table that takes the width of a route and returns the user-specified cost for that width. A user may enter any arbitrary width cost, but by default, PLEASE uses the costs shown in Table 3.3. To be consistent with the other heuristics, cost should be specified as a value between 0 and 1. Although corridor width is a real number and can have an infinite number of values, the getCorridorWeight discretizes width into three categories, small, medium, and large. This is done to eliminate meaningless differences between corridors. Two corridors that differ only slightly in width are equally likely to indicate a main route and should be treated equally. The cutoff values for these categories can be set by the user so that the values are appropriate for the current building layout. By default, PLEASE uses the values shown in Table 3.4. Agents in the simulation measure width at the entry point of the corridor.

$$W(x) = w_w * getCorridorWeight(x) \tag{3.5}$$

The congestion cost of route x, $Cong(x)$, is given by Formula 3.6, where w_{cg} is the weight applied to the congestion heuristic, sp_i is the desired speed (a normally distributed

Table 3.3: Width Costs Used by Default in PLEASE

Width	Cost
Small	1.0
Medium	0.5
Large	0.0

Table 3.4: Width Discretization Cutoff Values Used by Default in PLEASE

Discretization	Cutoff
Small	1.5 m
Medium	2.5 m
Large	>Medium

parameter value unique to each pedestrian) of pedestrian i, sp_j is the speed of pedestrian j, s_1 is 1 if $sp_j < sp_i$ or 0 otherwise, $n_{p,x}$ is the number of pedestrians along route x, and n_p is the total number of pedestrians. This formula assigns cost based upon the desired speed of the pedestrian and the current speed of pedestrians along the selected route. For each pedestrian along the selected route, if his or her speed is slower than the desired speed, then a cost is incurred relative to the speed difference. The cost is raised to the square so that smaller speed differences count less than larger differences. Finally, the result is normalized by the worst case cost (i.e., if every pedestrian in the simulation was along the selected route and was not moving).

$$Cong(x) = w_{cg} * \frac{\sum_{j=0}^{n_{p,x}} ((sp_i - sp_j) * s_1)^2}{sp_i * n_p} \tag{3.6}$$

The consensus cost of route x, $C(x)$, is given by Formula 3.7, where w_c is the weight applied to the consensus heuristic, $a_{j,x}$ is the angle between pedestrian j's orientation and the direction to route goal x from agent j (see Figure 3.2), and $n_{p,i}$ is the number of pedestrians surrounding pedestrian i.

$$C(x) = w_c * \frac{\sum_{j=0}^{n_{p,i}} (a_{j,x}/\pi)}{n_{p,i}} \tag{3.7}$$

The previously visited cost of x, $V(x)$, is given by Formula 3.8, where w_v is the weight applied to the visited heuristic, and $visitedCount(x)$ is the number of remembered times route x has been visited by the pedestrian in this simulation run. Each pedestrian is capable of remembering a specified limit of number of routes for a specified amount of time to reflect the finite memory of pedestrians. These limits may be set to any arbitrary value by the user, but by default, PLEASE has a limit of 10 routes for 1000 seconds.

$$V(x) = w_v * visitedCount(x)/maxMemory \tag{3.8}$$

3.5 Experimental Design

To test the effectiveness of the various heuristics described above, we use a combination of actual building layouts and building layouts constructed for the purpose of these experiments. In the buildings shown in Figure 3.3, blue lines represent walls of the buildings, dashed green lines represent exits signs, solid red lines represent emergency exit signs, and wavy gold lines represent door signs. The USU Business building (see Figure 3.3(a)) is an approximation of the ground floor of the actual building found on the Utah State University Campus. Likewise, the CSULB FM building (see Figure 3.3(c)) and the CSULB UP building (see Figure 3.3(b)) are approximations of the actual buildings found on the California State University, Long Beach campus.

The USU Business building and the CSULB UP building are used because the floor plans follow expected conventions: rooms do not function as hallways, main areas have wide corridors, and other accepted conventions are followed. The CSULB FM building is used because of the lack of convention with room placement. Rooms are found within rooms and several rooms can even function as hallways. The custom building is designed to represent a standard building. The corridor widths are not an exact indication of main routes and areas, but are still closely correlated to main routes. Rooms do not function as hallways.

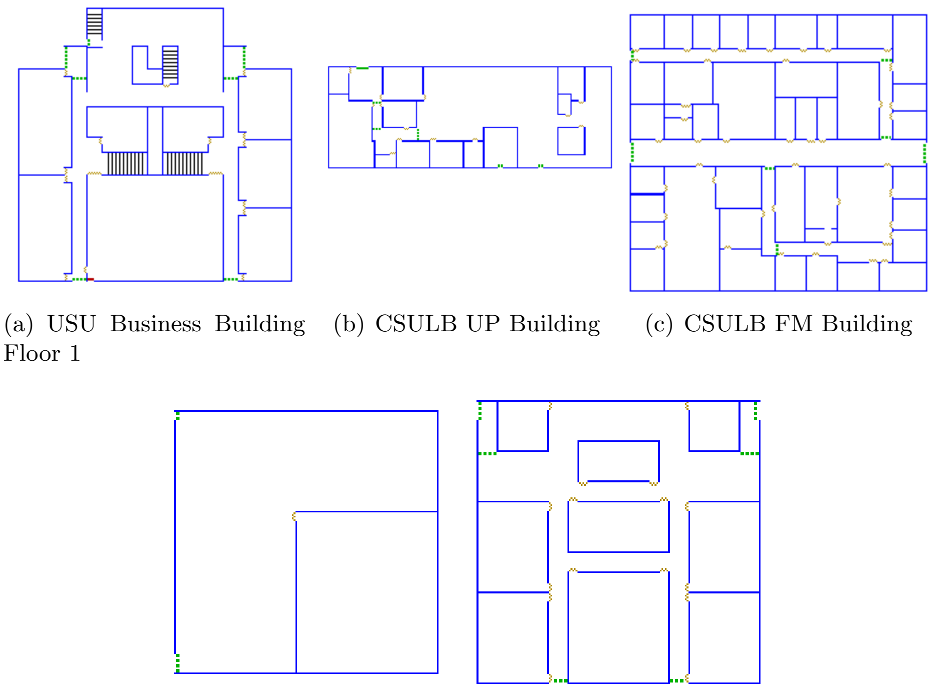

(a) USU Business Building Floor 1 (b) CSULB UP Building (c) CSULB FM Building

(d) Simple Building (e) Custom Building

Figure 3.3: Building layouts used in the heuristic evaluation experiments.

The number of rooms and routes are not so many as to be completely unmanageable by pedestrians either. Finally, the simple building is designed for ease of navigation. Navigating this building should be relatively simple, so heuristics that lead to prolonged egress times in this building are indicative of a problem.

For each building, we conduct a variety of tests. We measure the total egress time of 100 pedestrians, randomly distributed throughout the rooms and averaged over 20 simulations, using a single active heuristic. In all the tests, the previously visited heuristic is always active and is given a weight of five. This represents the agents' unwillingness to backtrack along a route previously traveled.

The viability of the signage heuristic is completely dependent upon the type and amount of signage found with the buildings. For this reason, a comparison between the signage and simple signage heuristics is helpful. The simple signage heuristic represents the same building stripped of all signs except for exit signs at actual exit doorways.

Based upon the performance of the individual heuristics, we then combine multiple heuristics, weighting each heuristic by its relative performance in relation to the total egress time achieved by the heuristics. These weights are further adjusted to improve the performance of the combined heuristics. The weight values for these experiments are shown in Table 3.5. Determining the exact weight specification to optimize performance when multiple heuristics are used is outside of the scope of this paper. However, the weights that are used provide good performance. Values of the parameters are easily modified.

Table 3.5: The Heuristic Weight Values Used When Combining Multiple Heuristics.

Heuristic	Weight
Shortest Leg	0.5
Simple Signage	1.0
Width	0.5
Congestion	3.0

3.6 Results and Analysis

In this section, we discuss the results of the experiment in four parts. First, we look at the effect of building layouts on egress times. Second, we consider in detail the result of applying only one heuristic at a time. Third, we look at the results of applying multiple heuristics simultaneously. Fourth, we rank each heuristic using a relative performance ratio.

3.6.1 Building Layouts

The USU Business building and the CSULB UP building have similar results for most of the heuristics (see Figure 3.4). This makes sense because both buildings have similar design characteristics. The main differences in heuristic performance occurred with the shortest-leg heuristic. The USU Business building layout happens to be conducive to using the shortest-leg heuristic. The outside rooms can actually be used as hallways that lead directly to the exits, and this is exactly what happens when the shortest-leg heuristic is used. The ground floor of the USU Business building layout appears similar to the custom building layout, yet the slight differences in spacing and room layout create large differences in egress times for the two buildings.

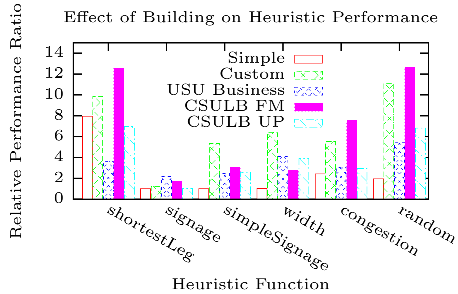

Figure 3.4: Effect of building layout on heuristic effectiveness. Performance ratio is calculated as the average time taken for 90% of the pedestrians to evacuate using the given heuristic divided by the average time taken for 90% of the pedestrians to evacuate using perfect knowledge.

The CSULB FM building and the custom building both have similar results in egress times for the shortest-leg, signage, congestion, and random heuristics. Although the actual layouts of the building are quite different, the underlying patterns are similar: width correlates with egress routes, long hallways have many adjoining rooms, and exits are distributed in a uniform manner. The main difference between the results for each building is the performance of the width heuristic. In the CSULB FM building, corridor width corresponds closely with egress routes, while in the custom building the correlation is weaker. The simple signage heuristic also yields different results in these two buildings. In the custom building, the top left and top right exits are not visible throughout most of the building and are thus highly underutilized. Meanwhile, the exits in the CSULB FM building have a much higher level of visibility throughout the building and are used more effectively.

The simple building also has different egress times compared to the other buildings. This building layout is simple and has few route choices. Thus, which route selection heuristic is applied makes little difference. The exception to this is the distance heuristic which causes extreme congestion around the inner room doorway. A large percentage of the pedestrians are closest to that route as the inner doorway is situated near the center of the building, making it the closest visible route to the majority of pedestrians in the simulation.

3.6.2 One Heuristic

When only a single heuristic is used, the results vary greatly between heuristics (see Figures 3.5(a)-3.5(e)). The signage heuristic performs well in all types of building layouts tested. This suggests that if pedestrians choose an egress route based upon well-designed signage, pedestrians can efficiently egress from a building even when completely unfamiliar with the layout of the building. Using simple signage, the egress times are still as good as or better than any other heuristic in the building layouts considered. This highlights the importance of even minimal building signage in assisting in pedestrian egress.

The shortest-leg heuristic leads to slow egress times in almost every building layout considered. In many instances, the shortest-leg heuristic does not even outperform a random choice policy. As discussed in Section 3.3, when the end goal is not known, choosing the route that is closest to the pedestrian becomes the shortest-leg first heuristic. This greedy route selection heuristic provides no guarantee that the route chosen will even lead to a direct exit. Additionally, when distance is the sole means of evaluating a route, congestion is a common occurrence. Pedestrians who are closest to a given doorway will select that doorway regardless of what side they are on or which direction other pedestrians are moving. Thus, the pedestrians on opposite sides of the doorway will converge at the doorway causing a bottleneck, and pedestrian flow rates through the doorway will be greatly inhibited. This is exactly what happens in the simple building and is the reason for the extremely delayed egress times (see Figure 3.5(e)). Interestingly, the shortest-leg heuristic performs remarkably well in the USU Business Building (see Figure 3.5(a)). This building is configured ideally so that greedily selecting the closest visible route actually leads pedestrians to an exit in a fairly efficient manner. One reason that this is the case is the double doors on most of the rooms. This allows the pedestrian to explore routes without having to backtrack, which is discouraged by the algorithm. Additionally, the end rooms have doorways adjacent to exits, which facilitates egress in this situation.

Similar to the shortest-leg heuristic, the fewest turns heuristic also leads to poor performance. The results are not shown here. The intuition behind the fewest turns heuristic is

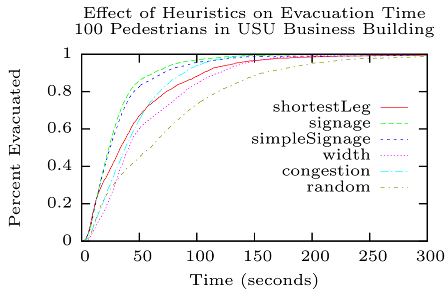

(a) USU business building heuristics results

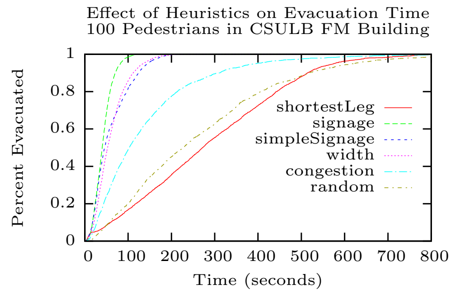

(b) CSULB FM building heuristics results

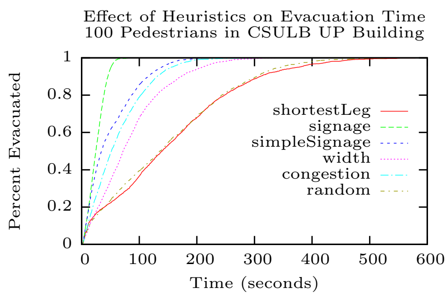

(c) CSULB UP building heuristics results

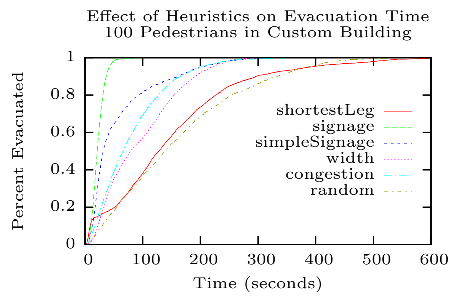

(d) Custom building heuristics results

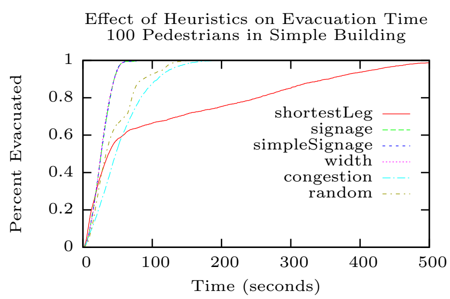

(e) Simple building heuristics results

Figure 3.5: Resulting egress times when agents use a single heuristic function to select an egress route.

to select a route that is as direct as possible. However, considering only the next route goal is too short-sighted and leads to routes that are drastically less direct than they could be. Without prior knowledge about the building layout, though, this short-sightedness cannot be overcome.

For the building layouts considered in this paper, the width heuristic leads to average egress times when compared to the other heuristics. For the simple building, the width heuristic actually leads to excellent egress times as the widest route is also the best exit. For the CSULB-FM building, choosing the widest route leads to finding an exit sooner than selecting a route by any other heuristic except for signage. In the remaining buildings, the width heuristic performs worse than the congestion heuristic, but still significantly outperforms a random policy.

The congestion heuristic does not necessarily provide an indication of which route leads to an exit, especially when none of the pedestrians have any knowledge regarding the building layout. However, avoiding congestion still improves the overall egress time by helping prevent bottlenecks and increasing the overall smoothness of pedestrian flow. This allows more routes to be explored in less time, which leads to better egress times. Although (due to space limitations) the results are not shown here, the consensus heuristic also relieves congestions at bottlenecks and improves pedestrian flow so that routes can be explored in a more efficient manner. The consensus heuristic would most likely prove to be more valuable if some pedestrians have additional knowledge about the building layout. If other pedestrians imitate the route selection behavior of their peers using the consensus heuristic, not only would the smoothness of pedestrian flow increase, but also the pedestrians would be able to take advantage of the additional knowledge other pedestrians might have. To be most effective, the congestion and consensus heuristics should be combined with another heuristic such as width or signage both of which provide an indication of an egress route.

3.6.3 Multiple Heuristics

After considering each heuristic individually, we next combine several heuristics to further improve performance (see Figure 3.6). Although many different combinations could be tried, in this paper, we consider combining the simple signage and width heuristics (SS-W) and the shortest-leg, simple signage, width, and congestion (S-SS-W-C) heuristics. When width is the sole heuristic applied, the egress times are too slow to be reliable in an emergency. The simple signage heuristic is also slower than desired but is still the best alternative to the signage heuristic (which may not be realistic for many buildings) when

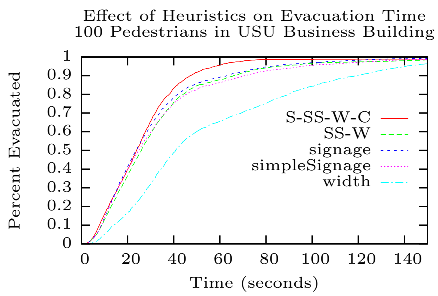

(a) USU business building combined heuristics results

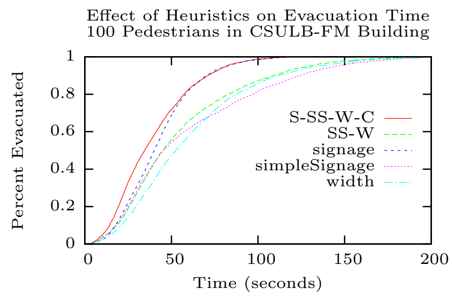

(b) CSULB FM building combined heuristics results

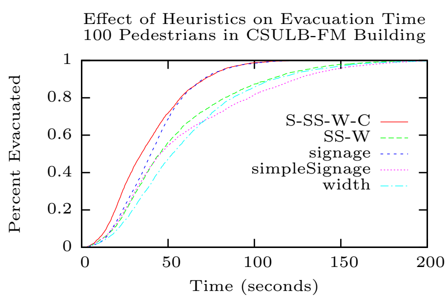

(c) CSULB UP building combined heuristics results

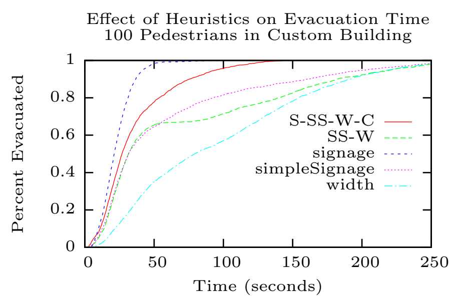

(d) Custom building combined heuristics results

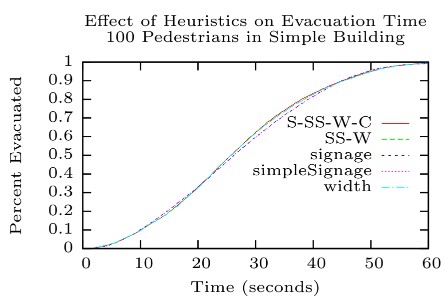

(e) Simple building combined heuristics results

Figure 3.6: Resulting egress times when agents combine multiple heuristics values into one heuristic function.

only a single heuristic is applied. The goal of combining width and signage is to take advantage of signage when available and to fall back on the width heuristic when signage is not available. We then include the other heuristics, namely shortest-leg and congestion, to further improve the egress times. For comparison purposes, the egress times of the signage, simple signage, and width heuristics are included in the charts.

As can be seen in Figure 3.6, combining the simple signage and width heuristic (SS-W) did indeed improve performance in most buildings in comparison to either heuristic alone. The custom building layout is the one exception. In this building, falling back on width proved to be detrimental to the overall egress time as the widest areas did not have direct exits. However, when several heuristics (shortest-leg, simple signage, width, and congestion, denoted as S-SS-W-C) are combined, performance is improved in every single building layout when compared singly to the performance of the individual heuristics. In most cases, the egress times matched or beat selecting routes based upon perfect signage. This indicates, that for a wide variety of buildings, the same heuristic functions can be applied to successfully egress from the building in a reasonably efficient manner using only the information that is directly perceivable by the pedestrian.

Although the shortest-leg heuristic does not perform well by itself, when combined with other heuristics, it leads to improved performance (results omitted due to space limitations). This is indicative of the value distance can play in route selection and justifies its use by actual pedestrians. However, it is important to note the disastrous impact that relying only upon distance can have upon the total egress time of individual pedestrians when no additional information is utilized.

3.6.4 Heuristic Rankings

The heuristic functions are ranked according to the relative performance ratio (RPR) of each heuristic in the above mentioned building layouts. For each building, the average time (t) it takes for 90% of the pedestrians to evacuate when each pedestrian has perfect knowledge of route distances and congestion levels is recorded. The average time (h) it takes for 90% of the pedestrians to evacuate when each pedestrian is using the heuristic function of interest is also recorded. The RPR of each heuristic function is then computed as h/t. Table 3.6 displays the average relative performance ratio for each heuristic.

Table 3.6: Ranking of Heuristic Functions by Relative Performance Ratio (RPR). A Lower RPR Signifies Better Egress Times.

Heuristic	RPR
signage	1.4364796
S-SS-W-C	1.5929914
SS-W	2.7136588
simpleSignage	2.8858816
width	3.6199478
congestion	4.29597
consensus	4.645231
random	7.6032004
shortestLeg	8.1949452
angle	9.4109684

3.7 Future Work

While several directions of future work are possible, we propose the following list as key questions needing to be answered.

1. In large buildings, could a pedestrian familiar with heuristics that work well in one area of the building apply the same heuristics in another area of the building to successfully egress from the building?

2. In this work we combined four heuristics, the shortest-leg heuristic, the signage heuristic, the width heuristic and the congestion heuristic. What other combinations are effective?

3. Which heuristic functions and weightings provide the most realistic egress results for emergency and non-emergency situations?

4. All the building layouts we considered are single story floor plans. How does the addition of multiple stories affect the heuristic functions used? Which floors have exits? How does stairway width come into play? Do pedestrians move first toward the relative exit location and then to the correct floor, or do they move first to the correct floor and then to the exit location?

3.8 Conclusion

Using heuristic estimations in selecting an egress route is a natural and common process performed by pedestrians on a daily basis, yet most simulation models do not adequately address this fact. This paper highlights the importance of including heuristic costs in pedestrian simulators, especially those designed to model egress in an emergency situation. As Ozel indicates in [38], when in stressful situations and under time constraints, pedestrians will react by filtering and bolstering information (i.e., relying more upon heuristic estimations). This leads to decisions that are sub-optimal and, as shown, can have a significant impact on the total egress time of the simulation. In emergency simulation models, it is not sufficient to assume pedestrians will make the best or even good choices, the simulation model needs to consider the possibility that pedestrians are forced to make split-second decisions with little information.

Golledge and others have shown that pedestrians frequently use distance as the primary heuristic in route selection. Many pedestrian simulators also use distance as the primary factor when selecting a route for pedestrian egress. However, when the total distance is unknown, using a greedy strategy of selecting the closest route available is seen to produce poor results in many circumstances. If additional factors are included in the decision, the distance heuristic can help improve egress times, even when the total distance is not known.

Using four main heuristics (the shortest-leg heuristic, the simple signage heuristic, the width heuristic, and the congestion heuristic, each appropriately weighted) is shown to produce good egress times even when no information about the building layout is known before the simulation begins. If only a single heuristic is used, the signage heuristic gives the best results even when the amount of building signage is minimal.

CHAPTER 4
SIMULATING KNOWLEDGE AND INFORMATION IN PEDESTRIAN EGRESS[1]

4.1 Introduction

In recent years, accurate pedestrian simulation has become an important research topic [19, 23, 39, 48]. Pedestrian simulation models can be employed in the design of safe facilities, validation of fire codes, and the automatic tracking and surveillance of pedestrians in live video feeds [1]. An important area affecting the accurate simulation of pedestrians is the route/path selection algorithm. Many different techniques are used, such as assuming perfect knowledge of egress routes and applying standard search and planning techniques, or generating a distance (cost) lookup table for each route choice. Other methods assume no prior knowledge of the possible routes or the building layout. Pedestrians can then only choose from those routes that can be locally observed by the pedestrian. While common, these methods are not an accurate representation of actual practice. Rarely would a pedestrian have complete knowledge of a building. Yet, having no prior knowledge about the building is also unrealistic for most cases. A few simulators allow route knowledge to be entered manually by the user to simulate different route knowledge for different pedestrians [15, 41], which may require a large time commitment by the user to properly set up the environment. In this paper, we propose a novel application of reinforcement learning to provide pedestrians with individualized knowledge of the building without requiring a large time commitment from the user. Pedestrians can learn about the environment in an initial learning phase, and then the actual simulation can be run with different pedestrians

[1] Feuz, Kyle and Allan, Vicki "Simulating Knowledge and Information in Pedestrian Egress" Condensed version accepted for publication in the 4^{th} International Conference on Agents and Artificial Intelligence 2012

having learned various routes.

Another factor that can affect route selection is congestion. The use of reinforcement learning to supply pedestrian agents with prior knowledge about the building can also be extended to include additional knowledge such as the average congestion levels of the different routes. Using this technique, we can analyze the effect that utilizing congestion knowledge has upon the egress time and efficiency of the simulation. In traffic management, studies conflict as to whether or not providing dynamic information about traffic congestion conditions improves the efficiency of the road network. Some studies indicate that providing such information can lead to road usage oscillation patterns as drivers switch between two alternate routes [52, 53]. To our knowledge, the question of the effectiveness of providing congestion information has yet to be answered regarding pedestrian egress. The effect of learning typical congestions levels in a building prior to the actual simulation is also unanswered. In this paper, we seek to fill these gaps by analyzing the effect of incorporating dynamic route congestion information and learned route congestion information into the route selection algorithm.

4.2 Related Work

Reinforcement learning has been studied extensively for several decades [26]. Different algorithms and techniques have been developed, each with benefits and drawbacks. In general, reinforcement learning algorithms can be divided into two broad categories: model-free learning and model-based learning. The main difference between these two techniques is that in model-based learning, an agent learns about both the transition relationship between states and the reward function, whereas in a model-free technique, an agent only learns about the reward function. For a good survey of reinforcement learning algorithms, consult [26].

Less studied, but receiving more attention recently, is reinforcement learning in multi-agent systems. One of the most influential works in multi-agent reinforcement learning is [31] by Littman. He extended the q-learning algorithm to adversarial zero-sum Markov games. Since this initial work, other algorithms for learning in cooperative environments have also been developed [6,32]. One of the challenges with multi-agent learning arises from the moving target problem [5]. Because each agent is learning simultaneously (and decisions by one agent affect desirable choices by others), the optimal policy may change over time. In addition to new algorithms, in some situations, multi-agent learning can ignore the presence of other rational agents and employ a traditional single-agent learning technique.

The effect of dynamic congestion information in traffic management is a well-studied topic that has not yet received much attention in pedestrian situations. Dia provides a framework for simulating driver behavior with dynamic route information. He leaves as an open question what effect such information will actually have upon route selection behavior [8]. Wahle et al. study the effect of dynamic congestion information in traffic scenarios [52, 53]. They use simulation models to predicate the effect that different congestion messages have upon traffic congestion. Their findings indicate that the results are dependent upon the type of information provided, but in general, dynamic information tends to decrease the overall network efficiency as oscillation patterns of road usage develop. Roughgarden shows that selfish routing does not minimize the total latency of a network and provides bounds upon the cost of selfish routing for several different latency functions and network topologies [45, 46]. However, using game-theory, Helbing et al. discover the emergence of alternating cooperation as a fair and system-optimal road usage behavior in a route choice game [20]. They conduct empirical tests using an iterated 2-4 player route choice game. Cooperation tends to emerge when individuals also exhibit exploratory behavior. They do not consider the case of providing dynamic information about the road conditions.

Although dynamic congestion information has not been heavily applied to pedestrian simulation, several researchers have included congestion consideration while modeling pedestrian egress. The work of Hoogendoorn and Bovy includes the cost of congestion when selecting routes and activities to perform [23]. The congestion information can either be derived from the pedestrian's current perceptions, or it can be based upon future predictions of congestion levels. How this information affects the overall efficiency of the system is not discussed. Banerjee et al. consider the complexity issues of dynamically discovering congestion and rerouting agents accordingly [2]. Their model assumes complete route distance information is known to pedestrians and that only pedestrians who perceive the congestion will choose new routes. This is in contrast to our model wherein route information may not be known and congestion may be known or estimated from previous experience even when the actual congestion cannot be directly perceived.

Pan represents one of the more comprehensive pedestrian behavioral models [39,40]. He includes pedestrian characteristics such as competitive, leader-following, altruistic, queuing, and herding. Using these characteristics, an agent considers visible routes before identifying its currently preferred route. Similarly, Koh defines an agent who only considers congestion and obstructions that can be directly perceived by the agent [27] . However, knowledge of the location of the end goals appears to be available to all pedestrians. A common simulation environment, buildingEXODUS, assumes the agents know about a set of user specified routes or all routes in the absence of the specification [15]. VISSIM, a commercially available pedestrian simulator, first processes the building layout to generate perfect route information for the pedestrians [41]. In VISSIM, the user also has the option of specifying specific routes for specific pedestrian sets [41].

4.3 Simulation Environment

Our study is performed using the Pedestrian Leadership and Egress Assistance Simulation Environment (PLEASE). PLEASE is built upon the multi-agent modeling paradigm wherein each pedestrian is represented as an individually rational agent capable of perceiving the environment and reacting to it. In PLEASE, pedestrian agents can perceive obstacles, hazards, routes, and other agents. The agents are capable of basic communication to allow for the formation and dissolution of coalitions and the sharing of knowledge. The agents use a two-tier navigational module to control their movement within the simulation environment. The high-level tier evaluates available routes and selects a destination goal. The low-level tier, based on the social force model [19], performs basic navigation and collision avoidance.

4.3.1 Knowledge Representation through Reinforcement Learning

Typically, reinforcement learning algorithms are used to discover a near-optimal policy. In fact, many reinforcement learning algorithms provably converge to the optimal policy [26]. One benefit of reinforcement learning to our simulation is the fact that when the search is truncated, a less than perfect solution is found. These solutions can be used to automatically generate various levels of pedestrian knowledge about the building configuration. These sub-optimal policies do have a unique constraint though as well: they must still be realistic. A learned policy that (when followed) never results in the successful egress of the agent is unacceptable. For this reason, we have implemented the reinforcement learning algorithm using model-based techniques. The details of the implementation follow.

Each agent builds a model of the building layout and the associated costs of available routes. To do this, the agent abstracts the building layout into a graph-based view. A common abstraction of building layouts is to represent rooms as nodes in the graph and doorways between rooms as edges in the graph. For the purpose of reinforcement learning, however, this abstraction is too course-grained. An agent is forced to associate a single cost (the edge weight) between two arbitrary, connected rooms. The true cost actually varies significantly depending upon the agent's location in the room. If time and space considerations are ignored, the building can be discretized into arbitrarily small grid cells, which allows the cost between nodes to be represented more accurately. Of course, this method is too costly in terms of time and space to be practical for buildings of even modest size. PLEASE uses a building representation in between these two extremes. To do this, we introduce the concept of decision points. A decision point is simply a point in the building at which an agent must decide in which direction he will proceed. These points may be placed at any arbitrary location, but in our models, the decision points are always placed at doorways and corridor intersections. We select these locations because they are areas that pedestrians must pass through to move from one area of the building to another. This prevents the systems from forcing a particular path upon an agent. The nodes in the graph represent decision points in the building, and weighted edges between nodes represent the average cost of a path between two decision points. This provides more fine-grained control over the costs learned while still being manageable for larger buildings.

Pseudo-code for the learning algorithm is shown in Figure 4.1. Initially, the agent's model is empty as the agent has no prior knowledge about the building. Each time an agent passes through a decision point, the agent estimates the cost (based upon distance and/or congestion levels) to all other visible decision points in the room. (See Formula 4.1-4.3). Additionally, the agent estimates the cost to other decision points known by the agent to be in the room. The weighted edge between decision points is then updated to reflect the newly estimated costs. Decision points that are not currently represented in the graph are added as necessary.

Definitions:
model - the adjacency matrix for the building layout representation
d - the decision point whose cost is being updated
dp - decision points in the same room as d
alpha - learning parameter of the algorithm, determines the weight applied to new cost estimates
estimateCost - estimates the cost between two decision points. See Formula 4.1 - 4.3
model.insert - inserts new rows and columns into the adjacency matrix as needed

```
Begin UpdateCost(DecisionPoint d)
foreach DecisionPoint dp in room
  if dp isVisible or isKnown
    cost = estimateCost(d, dp)
      if d, dp in model
        tmp = alpha * (cost - model[d][dp])
        model[d][dp] += tmp
        model[dp][d] += tmp
      else
        model.insert(d,dp,cost)
End
```

Figure 4.1: Algorithm used by the learning agent to update the estimated cost between decision points.

The agents estimate the cost from one point in the building ($dp1$) to another point in the building ($dp2$) based upon the distance and congestion levels between the two points. This estimate is specified by Formulas 4.1-4.3, wherein *cost* is the estimated cost of moving from $dp1$ to $dp2$, w_{cg} is the user-specified weight for congestion costs, w_d is the user-specified weight for distance costs, sp_i is the desired speed of the current pedestrian i, sp_j is the current speed of agent j, $n_{dp1,dp2}$ is the number of agents along the path from $dp1$ to $dp2$, N is the total number of agents, s_1 is 1 if $sp_j < sp_i$ and 0 otherwise, $dist(dp1, dp2)$ is the distance between $dp1$ and $dp2$, and $maxDistance$ is the maximum distance between any two connected decision points which is defined as the length of the diagonal of the building.

Both the distance cost and the congestion cost are weighted by user-specified parameters so that different relative weights can be chosen. Agents in the simulation are able to accurately estimate the distance to visible points within the simulation model as well as being able to estimate the distance to points they have previously visited. The distance is normalized using the maximum distance between two points on the simulation map. The congestion cost is estimated using the difference in speeds between the current pedestrian and other pedestrians that exist between the two points in the building. For each pedestrian along the selected route, if its speed is slower than the desired speed of the pedestrian, a cost is incurred relative to the speed difference. The cost is raised to the square so that smaller speed differences count less than larger differences. Finally, the result is normalized by the worst-case cost (i.e., if every pedestrian in the simulation was along the selected route and was not moving).

$$DistCost = \frac{w_d * dist(dp1, dp2)}{(maxDistance)} \tag{4.1}$$

$$CongCost = w_{cg} * \frac{\sum_{j=0}^{n_{dp1,dp2}} ((sp_i - sp_j) * s_1)^2}{sp_i * N} \tag{4.2}$$

$$cost = CongCost + DistCost \tag{4.3}$$

At this point, the agent must select the next route to follow. Pseudo-code for the route selection algorithm is shown in Figure 4.2. To do this, the agent performs a breadth-first search starting from each known decision point (dp) in the current room. If a path is found from the decision point to an end goal (g), the cost of the path is computed as the cost to dp plus the learned cost from dp to g. If no path is found to g, the cost is computed as the cost of d plus $UNEXPLORED\_COST$. $UNEXPLORED\_COST$ is a user-specified parameter representing the cost of choosing a route whose destination is not known. With probability p, the agent selects a random decision point to proceed towards, and with a probability of $1 - p$, the agent selects the decision point of least cost. The probability factor represents the probability an agent chooses to explore a different route. When the agent is learning, we set this probability to 0.15. This value reflects the speed with which agents learn a building. When learning congestion cost, this value also affects the reliability of the learned congestion costs. Agent training happens concurrently for all agents in the simulation. This creates a moving-target problem because congestion levels are constantly fluctuating as agents change their respective policies based upon the congestion levels encountered previously. When the probability of exploring is high, a large number of agents will not use routes they normally would if they were not exploring, thus leading to inaccurately learned congestion costs.

Definitions:
dp - decision point in the current room to consideration
$explore$ - normally distributed random value between 0-1
p - probability of exploration
$cost$ - dictionary of costs for decision points considered
$estimateCostTo(dp)$ - similar to estimateCost in Figure 4.1 but uses the agents position as $dp1$
$BFSCost(dp)$ - the cost found by performing a breadth-first search from dp to the end goal

```
Begin routeSelection()
  foreach DecisionPoint dp in room
    cost[dp] = estimateCostTo(dp) + BFSCost(dp)
  explore  random()
  if explore <= p
    return random DecisionPoint in room
  else
    return arg min cost[dp]
End
```

Figure 4.2: Algorithm used by the agent to select the desired route of travel

4.4 Congestion Considerations

As we are interested in the effects of congestion on the egress efficiency of the system, we consider the two cases: 1) ignore current congestion levels and 2) adjust decisions based on directly perceived congestion.

4.4.1 Ignore Congestion

We use the case of ignoring congestion as a base case against which we can compare all other cases. For many situations, we expect that completely ignoring congestion will lead to slower egress times as the building corridors are used inefficiently. However, ignoring congestion is still a feasible pedestrian behavior. Generally, pedestrians prefer to travel along paths they have previously traveled [38]. This may mean that, in spite of congestion, they continue to travel along their preferred route. Congestion might also be ignored if the pedestrian believes that other routes will not decrease their egress time.

4.4.2 Adjust to Directly Perceived Congestion

Adjusting to directly perceivable congestion is common in many simulation models [23, 27]. Intuitively, it makes sense that pedestrians adjust their route based upon perceived congestion. From a modeling perspective, this case has the additional benefit of not requiring any additional knowledge about congestion in other areas of the building. The question remaining is this: Does it improve the overall egress times?

4.5 Knowledge Considerations

We are interested not only in the effect of reacting to congestion upon egress times, but also in the effect congestion knowledge has upon egress times. We consider three types of knowledge that pedestrians may have: 1) learned route distance knowledge 2) learned route congestion/distance knowledge, 3) system-provide route congestion/distance knowledge.

4.5.1 Distance Knowledge

The case of route distance knowledge represents pedestrians who have learned route distances but not route congestion levels. These pedestrians' primary concern is arriving at the destination rather than the congestion levels along the way. The completeness of the distance knowledge a pedestrian has is dependent upon the amount of training the agent has. In this paper, we consider agents with both low knowledge (10 training runs) and high knowledge (100 training runs).

4.5.2 Congestion Knowledge

Knowledge of the average congestion costs is more reflective of reality as pedestrians familiar with a building are also typically familiar with the route usage patterns. This case assumes that pedestrians remember congestion costs from previous experience in the building in addition to the distances between various decision points. In this case, the congestion costs are associated with routes between decision points. Each time an agent travels a given route, the expected cost for that route is then updated. The completeness of the distance/congestion knowledge that a pedestrian has is dependent upon the amount of training the agent has. In this paper, we consider agents with both low knowledge (10 training runs) and high knowledge (100 training runs).

4.5.3 System Knowledge

The final case we consider is providing pedestrians with dynamic route congestion information and route distance information. This allows a pedestrian to evaluate all possible routes for distance and congestion, even when those routes are not directly perceivable (i.e., the route cannot be seen). Such information may one day be generally available to pedestrians through personal hand-held devices or public displays [3,28,36]. Simulating the knowledge provided by these egress assistance devices can provide an early look into the effectiveness such devices may have on pedestrian egress times.

4.6 Experimental Setup

All the experiments conducted in this paper use four different building layouts (see Figure 4.3). Building A is designed with specific congestion considerations in mind. To

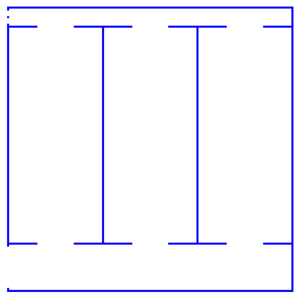

(a) Building A

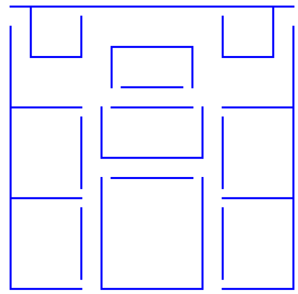

(b) Building B

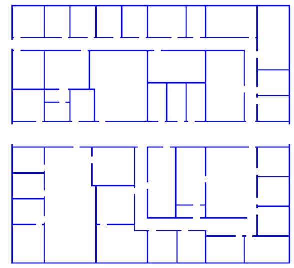

(c) Building C

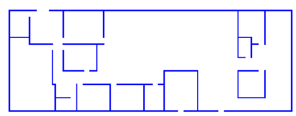

(d) Building D

Figure 4.3: Building layouts used in the congestion experiments.

pedestrians in the inner rooms, each room doorway appears to be of equal value. However, the lower doorways lead to a wider corridor and exit and will thus be able to accommodate more pedestrians. Building B is designed to be representative of a general building layout. Buildings C and D are approximations of actual buildings found on the California State University, Long Beach campus.

4.6.1 Experiment 1

The purpose of the first set of experiments is to demonstrate the feasibility of using reinforcement learning as a means to represent pedestrian knowledge in a simulation environment. To do this, we show that as the number of learning trials (to which an agent is subjected to) increases, the amount of building knowledge the pedestrian acquires also increases. We show that this increase in building knowledge leads to a corresponding decrease in pedestrian egress times.

For each building, we conduct the test as follows. Five-hundred pedestrians are trained in the building for 100 simulation runs during which the agents learn route distance costs. After each simulation run, the agents' current policy is saved to disk so that we can recover the policy learned after any given number of simulations runs.

In order to determine how much knowledge an agent has gained about a particular building, we first need to define some metrics. We consider three key factors affecting route knowledge: 1) the number of known decision points (node knowledge), 2) the number of known paths between decision points (edge knowledge), and 3) the number of decision points known to be direct exits (exit knowledge). Using these metrics, we can then calculate the average amount of knowledge obtained by the agents for each trial run.

Figure 4.4 shows the average effect of multiple training runs on the total knowledge an agent has. As can be seen from the graphs, the different metrics indicate different knowledge levels, but the values of all metrics show an increase as the number of training runs increases. Agents quickly learn a high percentage of the decision points and paths between decision points, but for the key decision points representing building exits, the percentages are lower. This indicates that although agents learn many internal routes after 100 training runs, they are learning different exits at a slower rate.

As the amount of knowledge pedestrians have increases, so should the efficiency with which agents egress from the building. To test this, we measure the egress time of 500 pedestrians randomly distributed throughout the building, averaged over 20 simulations using policies of various training levels. Averaging the results over 20 simulation runs provides relatively small error bars which boost our confidence in the accuracy of the mean egress times obtained for each training level. Figure 4.5 compares the average egress times obtained when agents have gone through 10, 50, and 100 training runs for each building.

The effect of additional training in building A is minimal. This implies that the additional knowledge gained is not helpful in improving egress times. Building A is fairly simple and therefore the general layout can be learned quickly. Building B and building C both show substantial improvement in egress times as the number of training runs increases indicating that the knowledge gained by the pedestrians is indeed helpful in improving egress times. Building D shows substantial improvement in egress times between 10 and 50 training runs, but then little change occurs between 50 and 100 training runs. This correlates to

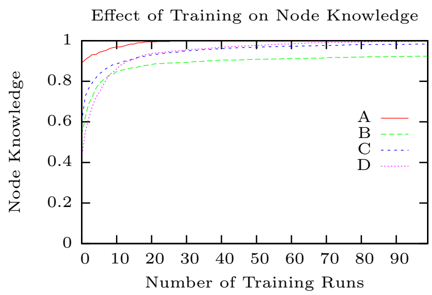

(a) Percentage of nodes learned over time

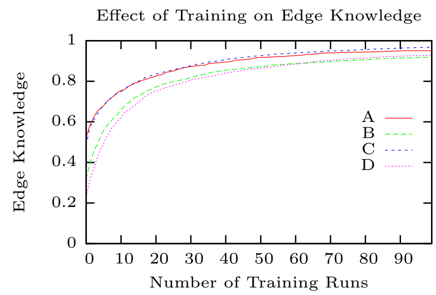

(b) Percentage of edges learned over time

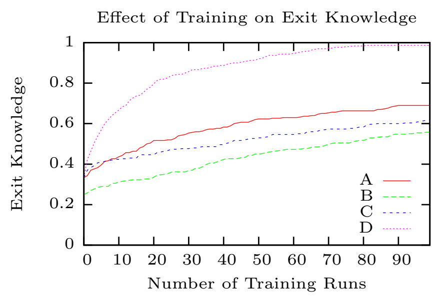

(c) Percentage of exits learned over time

Figure 4.4: Percentage of knowledge gained over time using three metrics

the previous results in Figure 4.4, where the amount of knowledge gained between 50 and 100 training runs is much less for building D than it is for the other buildings.

4.6.2 Experiment 2

The next set of experiments is intended to measure the effectiveness of learning average route congestion costs in addition to route distance costs. The experiments also measure the effectiveness of reacting to currently visible congestion and adjusting the selected route accordingly. Notice the distinction between learning congestion levels and reacting to current congestion levels. 'Choose to react' to congestion or 'ignore congestion' do not imply either a knowledge or a lack of knowledge of average congestion levels. Each is merely a decision of whether or not to include current congestion levels in the decision-making process. Conversely, having knowledge of points of congestion does not imply that the agent must react to current congestion levels, only that the agent will consider previous learned

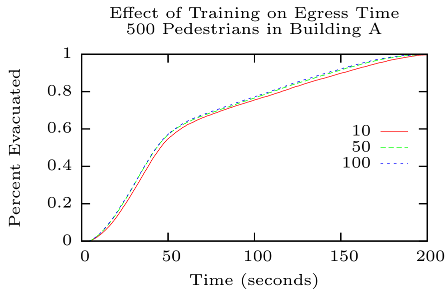

(a) Egress time for building A

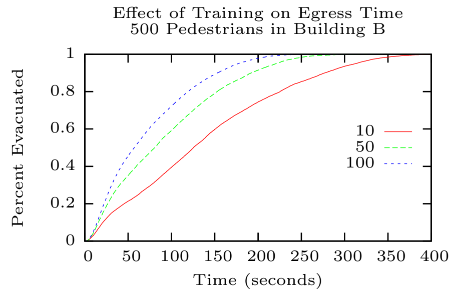

(b) Egress time for building B

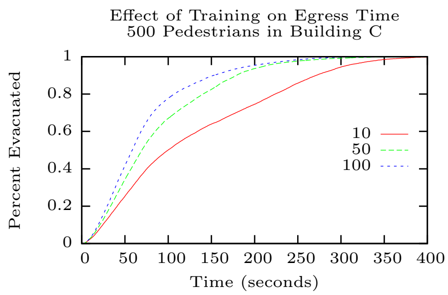

(c) Egress time for building C

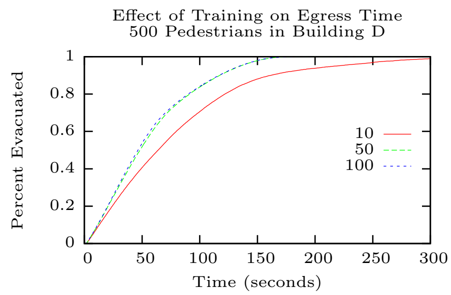

(d) Egress time for building D

Figure 4.5: Percentage of pedestrians exited over time using three levels of training

congestion levels when making the decision. Thus, an agent having no previous congestion knowledge can react to current congestion levels, and an agent having previous congestion knowledge can choose to ignore current congestion levels.

For each building, we conduct the test as follows. Five-hundred pedestrians are trained in the building for 100 simulation runs, learning both route distance costs and average congestion levels. The agents' current policies are check-pointed after 100 training runs so that we can compare the egress times when pedestrians have high levels of knowledge. We then measure the total egress time of 500 pedestrians randomly distributed throughout the rooms, averaged over 20 simulations. There are two parameters that we adjust in these tests: whether the pedestrian reacts to congestion, and what type of knowledge the pedestrian has. Pedestrians can either ignore current congestion levels or react to current congestion levels, and pedestrians can have either learned distance knowledge, learned congestion knowledge (which also includes distance knowledge), or system provided knowledge for both route distances and congestion levels. Therefore, we have six cases to consider: 1) ignore current congestion and have learned distance knowledge (Ign-Dist), 2) ignore current congestion and have learned both distance and congestion knowledge (Ign-Cong), 3) ignore current congestion and have perfect distance knowledge provided by the system (Ign-Sys), 4) adjust to congestion and have learned building distance knowledge (Adj-Dist), 5) adjust to congestion and have learned both distance and congestion knowledge (Adj-Cong), and 6) adjust to congestion and have perfect distance and congestion knowledge provided by the system (Adj-Sys).

The results are shown in Figure 4.6. In every building layout tested, agents who have learned both route distances and congestion levels have faster egress times than agents who have learned only route distances. This indicates that learning average congestion levels and using that knowledge in pedestrian egress is beneficial. However, the same cannot be said about reacting to congestion. Although learning average congestion levels is always beneficial in our tests cases, reacting to current congestion returns mixed results. In building A, reacting to current congestion always improves performance regardless of the type of knowledge a pedestrian has. This is not surprising because building A is specifically designed to contain severe congestion problems that are easily mitigated. In Buildings B, C and D, reacting to current congestion yields little change in overall egress time except when pedestrians have only route distance knowledge. In this case, reacting to the current congestion levels actually decreases the overall performance of the agents. This most likely occurs because the pedestrians are uniformly distributed within the building so the congestion is also well distributed. Thus, when a pedestrian chooses to take an alternate route, he or she soon discovers that it is equally congested. Finally, in building D, reacting to congestion actually improves performance if the pedestrian has learned previous congestion levels. Although the pedestrians are still uniformly distributed, the routes to the exits are not. In this case, knowing the typical congestion levels allows an agent to make a better decision when reacting to the current congestion levels.

Interesting patterns in the data can also be observed when the egress times of agents

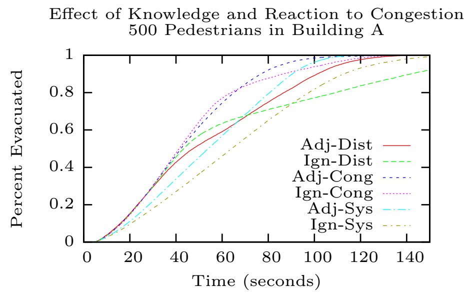

(a) Egress time for building A

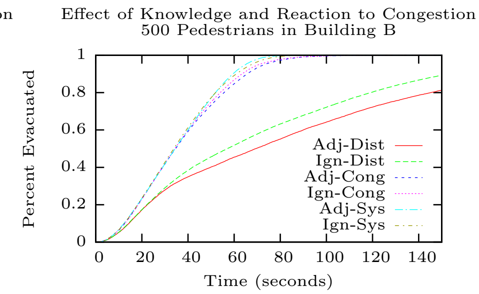

(b) Egress time for building B

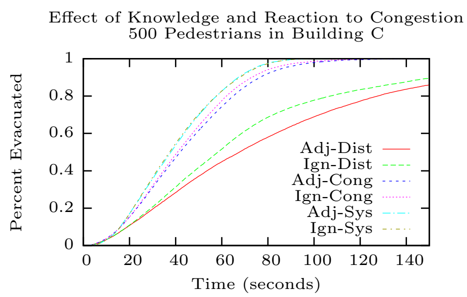

(c) Egress time for building C

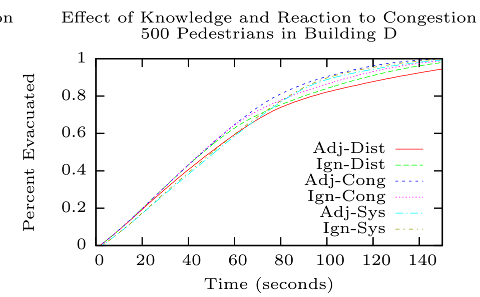

(d) Egress time for building D

Figure 4.6: Percentage of pedestrians exited over time with different amounts of knowledge and behaviors

with different types of knowledge are compared. In half the buildings (A, and D) utilizing learned congestion levels provides the best egress times, even outperforming system provided information. This is probably due to the oscillation that can occur when dynamic information is provided. As is also seen in traffic management, providing dynamic information can lead to many pedestrians switching routes simultaneously which decreases the efficiency with which pedestrians are able to evacuate the building. In every building layout tested, when pedestrians have only learned distance information, the performance is the worst of all possibilities considered. Interestingly though, a pedestrian having system information but ignoring current congestion levels and using only distance information is able to egress from most buildings quickly. However, the distance information of such a pedestrian is complete. One would expect that with enough training, pedestrians having learned only distance cost would also be able to egress from buildings with similar efficiency.

4.7 Future Work

The results presented in this paper were obtained using a limited number of building layouts. Would the results be consistent for a larger number of actual building layouts?

How do the egress times of different levels of training correspond with egress times using other knowledge representation schemes and how do they compare with reality? What level of knowledge is reasonable for typically scenarios?

Adjusting to congestion yields ambiguous results. Can we determine what factors and characteristics determine when pedestrians should react to congestion levels?

4.8 Conclusion

Providing agents with perfect knowledge is unrealistic for many pedestrian egress situations. However, manually specifying specific route knowledge can be a difficult and time-consuming task. We have shown that reinforcement learning can be applied to successfully represent different levels of knowledge about a building layout and produces egress times dependent upon the knowledge level of the pedestrians. We have also provided three different metrics for measuring the amount of building knowledge an agent has.

Using reinforcement learning, we have also shown that learning congestion cost in addition to distance costs leads to quicker egress times. However, reacting to current congestion levels has ambiguous results. This is consistent with similar studies in the traffic management domain. The layout of the building is found to have an impact on the strategy a pedestrian should use to minimize egress time.

CHAPTER 5
GROUP FORMATION AND KNOWLEDGE SHARING IN PEDESTRIAN SIMULATION[1]

5.1 Introduction

In coalition formation theory, a coalition will only form if the utility achieved by the agents in the coalition is greater than the utility each agent could achieve alone [49]. This assumption of individual rationality is typically made about agents in utility theory [47]. In pedestrian egress, pedestrians frequently travel in groups. From the above assumptions, the utility of pedestrians should be greater by joining a group than if they were to travel as individuals, yet most of the literature indicates that group formation has a negative effect on flow rates, average speed, and egress times [35, 44, 55, 56]. From this data, two logical conclusions can be drawn, given that individuals do form coalitions. First, the utility of pedestrians during egress is affected by more factors than just egress time. Such factors may include altruism or social influences. Second, other benefits are gained so that the overall egress time is not always negatively affected. Such benefits may include information sharing or stress reduction. In this paper, we primarily consider the effect of information sharing on egress times and show that, in certain situations, egress time can be improved through the formation of groups.

To understand the effects of group formation in pedestrian egress through simulation techniques, it is necessary to have a pedestrian simulation framework that is capable of modeling groups and group formation. In this paper, we therefore propose a novel group framework to be used in simulating pedestrian egress. This framework, like several other

[1]Feuz, Kyle and Allan, Vicki "Group Formation and Knowledge Sharing in Pedestrian Simulation" submitted to 11^{th} International Conference on Autonomous Agents and Multiagent Systems 2012

works [35, 43], uses a social cohesion force to simulate group behavior. However, our framework also includes new techniques to address the question of route selection and information sharing in a group situation.

5.2 Related Work

Several researchers have found negative effects on flow rates, and average speed when group movement is considered. Moussaid et al. conduct empirical studies to determine several group parameters including size, structure formation, and speed [35]. The average speed is found to decrease with increasing group size. Similarly, Qiu develops a framework for group modeling in [43] that predicts decreased flow rates for pedestrian groups. Initially, the simulation model predicts an increase in flow rate as group size increases. As group size continues to increase, however, the flow rates decrease. This makes sense intuitively because with two or three people traveling together, the interactions are harmonious and increase the smoothness of pedestrian flow. However, as the group size increases, the flow rate begins to decrease as groups spend a large amount of time waiting for the other members of the group to pass through a doorway or bottleneck before moving on. The model also predicts decreased flow rates as the strength of interaction between group members (intra group strength) and between groups (inter group strength) increases. Both Moussaid and Qiu use a similar idea of representing group interactions through the use of a social cohesion force that simulates group members' desire to maintain a close proximity to each other. Neither model, however, addresses the issues of route selection and information sharing in a group setting.

Ji and Gao consider the effect of multiple leaders with perfect evacuation route knowledge [42]. In their simulation model, they find that including more leaders increases the egress efficiency, in terms of total egress time, up to a certain saturation point, after which including more leaders decreases the egress efficiency [42]. This occurs because as the number of leaders increase, more pedestrians receive conflicting directions from multiple leaders which in turn hinders their ability to quickly egress from the building. Murakami et al. conduct a similar test using fire drills in a simulated model [37]. Leaders can instruct evacuees to either follow them or to take a certain route. They term these the "follow-me" and"follow-direction" methods, respectively. In the simulation model, the follow-direction method is implemented by leaders giving directions to nearby visible pedestrians, and the pedestrians then always follow the directions given. For the follow-me technique, leaders instruct nearby pedestrians to follow them. If pedestrians fall behind, the leader waits for them. The follow-me technique is most effective when sufficient numbers of leader are present so that the pedestrians can easily find someone to follow [37]. In the described experiments, the dedicated leaders are known beforehand and given additional information regarding which exits to take. We study the more general case where leaders are not known or given special training prior to the simulation.

Yang et al. use a simulation model to measure the effect of grouping upon egress time [55, 56]. They classify grouping as either spatial or directional [56]. Spatial grouping relates to individuals' desire to be close physically. Directional grouping relates to the desire of individuals to move in the same direction as others. Their model indicates that spatial grouping is detrimental to egress efficiency. However, directional grouping is found to increase the egress efficiency. This is one of the few papers that shows any benefit to grouping. However, we are interested in showing that even spatial grouping can lead to additional benefits.

Ozel applies the theories of decision-making under time pressure and stress to the emergency egress of pedestrians [38]. Under these theories, pedestrians will use different coping mechanisms to process the information available and make decisions. These coping mechanisms include increased rate of processing, avoidance of decision-making, and subjective filtration of the information. He indicates that joining a group can be viewed as a stress coping mechanism that allows a pedestrian to avoid the decision-making process by passing the responsibility of making the decision on to the leader of the group.

5.3 Our Model

Our research study is performed using the Pedestrian Leadership and Egress Assistance Simulation Environment (PLEASE), which we developed for this purpose. PLEASE is built upon the multi-agent modeling paradigm wherein each pedestrian is represented as an individually rational agent capable of perceiving the environment and reacting to it. In PLEASE, pedestrian agents can perceive obstacles, hazards, routes, and other agents. The agents are capable of basic communication to allow for the formation and dissolution of coalitions and the sharing of knowledge. The agents use a two-tier navigational module to control their movement within the simulation environment. The high-level tier evaluates available routes and selects a destination goal. The low-level tier, based on the social force model [19], performs basic navigation and collision avoidance.

5.3.1 Route Selection

PLEASE uses the concept of decision points to facilitate pedestrian route navigation. A *decision point* is defined as a point in the building at which an agent must decide upon the next location in the route. These points may be placed at arbitrary locations, but typically decision points are placed at doorways and intersections. When exiting from a building, pedestrians navigate from one decision point to another. By only placing decision points at doorways and intersections, the placement of the decision point does not require the pedestrian to pass through an area that they would not normally pass through when navigating from one area of the building to another.

Several different route selection algorithms are implemented in PLEASE. For this paper, we focus on two different route selection algorithms, a local route selection algorithm and a trained route selection algorithm which are explained below. We use these two route selection techniques to compare differences in egress times when pedestrians have different amounts of knowledge about the building. The local route selection algorithm does not require any prior knowledge of the building as it uses only locally observable information. The training algorithm allows pedestrians to know the route costs for decision points with which they are familiar. This knowledge can then potentially be shared with other pedestrians.

Local Route Selection

The local route selection algorithm estimates the cost of exiting via a given decision point based upon several locally observable characteristics of the point. In this paper, we use the distance, corridor width, room signs, and congestion characteristics when estimating costs. These characteristics have been found to work well for a variety of building layouts [9].

Trained Route Selection

The training algorithm allows agents to experience multiple simulation runs in a building during which time the agents may learn the expected costs to different decision points. This information is stored in the agent's model of the building. More training provides agents with more knowledge about the building, allowing for more effective route planning. In this study, we restrict the learning to distance information. This is done to facilitate the sharing of knowledge between group members by enforcing a common cost metric as discussed in Section 5.3.4.

5.3.2 Group Formation

PLEASE allows for pedestrians to travel in groups. These groups may be formed to share knowledge, request aid, relieve stress, or interact socially. In PLEASE, groups are more than just individuals traveling in the same direction. Being a member of a group implies communication, agreement, and a desire to remain close together. Pedestrian groups may be of a static or dynamic nature. All groups in PLEASE have a leader who is ultimately responsible for the group decisions such as where to go next or when to allow other pedestrians to join the group.

Similar to [35, 43], we represent groups using an additional cohesive force applied between group members. Pedestrians in a group seek to maintain a certain proximity to visible group members. This makes sense because one pedestrian cannot maintain a certain proximity to another pedestrian unless the location of the other pedestrian is known. When turning the corner, pedestrians might temporarily lose sight of one another, but as the other pedestrians also turn the corner they are able to reconnect.

Static Group Formation

Static groups can be formed at the beginning of the pedestrian simulation. These groups represent relationships that are defined outside of the simulation such as family, friend, or business relationships. Static groups do not change throughout the simulation: new members cannot be added and current members are only removed when they exit the building. Rather than require the user to define each group manually, PLEASE uses user-defined parameters to automatically create groups at the start of the simulation. Empirical studies have found that pedestrian group sizes tend to be small and follow a zero-truncated Poisson distribution [35]. This distribution can be approximated by adjusting the parameters controlling group formation.

Dynamic Group Formation

Dynamic groups can be formed throughout the simulation. An agent may seek to join or leave a group at any time during the simulation, but joining a group requires the consensus of the group members. Pedestrian agents use utility theory when deciding whether to join a group and whether to accept new group members. The two actions each have separate utility functions. We refer to the utility of joining a group as the agent's individual utility function. We refer to the utility of accepting new group members as the group utility function. PLEASE is built to be extensible, so the exact utility functions used may be easily changed. For dynamic group formation, we are interested in determining if there is a benefit in joining a group whose members have more knowledge of the building versus joining any group as a means to relieve stress. To make this comparison, we implement an individual utility function based on the current stress level of the agent, and a separate individual utility function based upon the amount knowledge the agent has of the building. Similarly, we define two group utility functions, one focusing on the stress level of the leader, and the other focusing on the knowledge level of the leader. We keep this function simple to allow for a more direct comparison between group formation based upon stress and group formation based upon knowledge. In describing the group formation process, we use the following notation.

- A - The set of agents in the simulation
- G - The set of groups of agents in the simulation.
- S - The set of agents in a group.
- S_x - The group of which agent x is a member.
- L_s - The leader of S

Initially $\forall x \in A$, $S_x = \{x\}$ and $L_s = x$. This states that at the start of the simulation, each agent in the simulation is the leader of a group consisting only of the individual agent. The group formation process is divided into three steps, request admission, extend invitation, accept invitation. Any agent may request admission into any nearby group. Any agent may accept any received invitation. Only group leaders may extend invitations to other agents. At every time step t, $\forall x \in A$, agent x evaluates its utility function and then requests admission and/or accepts invitations to nearby groups based upon the expected utility. $\forall S \in G$, L_s evaluates its group utility function and extends invitations based upon the expected utility.

5.3.3 Route Consensus

The PLEASE model allows groups to use three different route consensus mechanisms that incorporate suggestions from group members to different degrees. The route consensus mechanisms are least-cost route (LC), most-common route (MC), or dictator. The dictator mechanism simply chooses the decision point proposed by the group leader and serves as a benchmark to measure the effectiveness of groups. LC and MC require each pedestrian in the group to submit a preferred decision point and an associated route cost estimate. From the proposed routes, the LC mechanism then selects the route with the least cost (as identified by group members). In order for the LC mechanism to work successfully, group members must be using the same scale to measure cost. The group members are assumed to be reliable. The MC mechanism selects the most frequently proposed route. Ties are broken by average route cost. The MC mechanism mitigates the problems inherent in comparing costs computed by various means as it chooses the most commonly proposed route. The group members can have completely different route cost functions, and a minority of group members could be unreliable without affecting the route selected by the group. However, this mechanism is unable to take full advantage of the special knowledge that any particular agent may have.

5.3.4 Information Sharing

When using the training route selection algorithm, agents are allowed to share expected route cost information that they learn. Current simulation models assume that pedestrian knowledge is not shared among group members. This is a valid assumption for many situations. It represents pedestrian groups choosing an egress route without prior discussion as to which route is the most efficient or effective. However, pedestrian groups might also first discuss the benefits and drawbacks to a particular route before deciding on an egress route. PLEASE allows for either scenario, and in this paper, we consider the effects of both.

In PLEASE, expected route-cost information may either be public or private. If the information is public, then, at the beginning of the simulation, group members may share all the route cost information for each decision point learned during the training runs. Each agent has access to other group members' models to integrate into its own model. Currently, all model information is treated equally, so costs are integrated as an average of the other agents' costs. A more complex model might allow for issues such as trust and reliability to affect the weight that each agent applies to other agents' models while integrating the costs into their own model.

Sharing information requires communication costs. In an actual situation, sharing information may take anywhere from a few seconds to a few minutes. To account for this fact, the system user may specify a knowledge sharing cost, which is the time in seconds an agent spends sharing the route information with other group members. While route information is being shared, no member of the group moves towards any goal location. If group information is private, the group members do not share complete route-cost information. Hence, there is no associated communication cost.

As the route consensus techniques are combined with different information sharing techniques, understanding exactly what information is shared can be confusing. To help clarify, we state explicitly what information is shared for the various combinations. Using the LC consensus mechanism with public knowledge results in agents sharing all their respective knowledge for each decision point at the beginning of the simulation. Then, as the simulation proceeds, at every decision point, each agent proposes his or her preferred next decision point and an estimate of the total route cost via that decision point. The cheapest proposed decision point will be selected. Using the MC consensus mechanism with public knowledge results in agents sharing all their respective knowledge for each decision point at the beginning of the simulation. Then, as the simulation proceeds, at every decision point, each agent proposes his or her preferred next decision point and an estimate of the total route cost via that decision point. The most commonly proposed decision point will be selected. Because all the information has been shared previously, the only difference between these mechanisms is in the individual perspectives of the agents. One agent might have a clearer view of congestion than another agent, or one route might be closer to one agent but further away for another agent.

With the LC consensus mechanism and private knowledge, no knowledge is shared between agents at the beginning of the simulation. Then, as the simulation proceeds, at every decision point, each agent proposes his or her preferred next decision point and an estimate of the total route cost via that decision point. The cheapest proposed decision point will be selected. With the MC consensus mechanism and private knowledge, no knowledge is shared between agents at the beginning of the simulation. Then, as the simulation proceeds, at every decision point, each agent proposes his or her preferred next decision point and an estimate of the total route cost via that decision point. The most commonly proposed decision point will be selected. In this case, considerable differences exists between these mechanisms as each agent has unique knowledge.

5.4 Experimental Results

In this research, we consider the effects of static grouping on pedestrian egress times.

Various experiments control the route selection algorithm used, the route consensus mechanism, and the knowledge sharing available to pedestrians. Egress times are calculated with 100 agents per simulation. Test are repeated 20 times to put error bars into acceptable ranges. To quantify differences between the performances of the different mechanisms, we define efficiency as the amount of time taken to evacuate a given percentage of pedestrians. We use this definition of efficiency throughout our discussion of these experiments.

5.4.1 Static Groups

In the first experiment, we compare the results of static group formation when pedestrians use the heuristic route selection with and without group formation. The group consensus mechanism and knowledge sharing mechanism have little effect on the egress time, because in this test, none of the pedestrians have prior knowledge of the building and they all use the same heuristic function. The purpose of this experiment is to verify that pedestrian groups have a negative impact on egress time and to quantify that impact when no knowledge is shared between pedestrians. As can be seen in Figure 5.1, forming static pedestrian groups without sharing knowledge has a negative impact on egress times. Group formation is 29% less efficient at the 50% mark. In other words, the average time taken to evacuate 50% of the pedestrians is 29% greater when pedestrians form groups than when no groups are formed. Group formation is 35% less efficient at the 70% mark and 69% less efficient at the 90% mark.

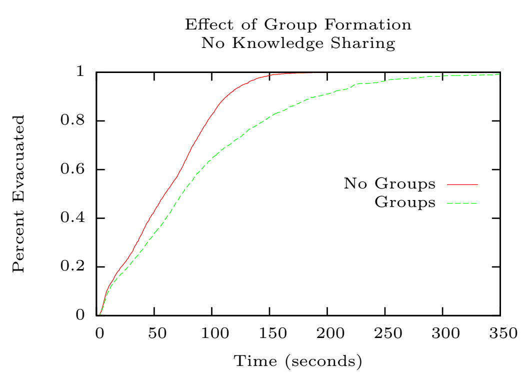

Figure 5.1: Comparison of static group formation on egress times when pedestrians have no prior knowledge of route costs. Group use the least-cost route consensus mechanism

In the second experiment, we show that group formation has a negative impact on egress times when pedestrians have individualized knowledge of the building (acquired through training) but do not share that knowledge. To do this, we compare the egress times of the pedestrians, which have learned route-distances over the period of 15 training runs, using no groups versus using a dictatorship.

Figure 5.2 shows the results of the second experiment. Again, as is expected, forming pedestrian groups leads to slower egress times compared to no groups. Specifically, the dictatorship group consensus mechanism leads to egress times that are 20% less efficient at the 50% mark, 32% less efficient at the 70% mark, and 69% less efficient at the 90% mark.

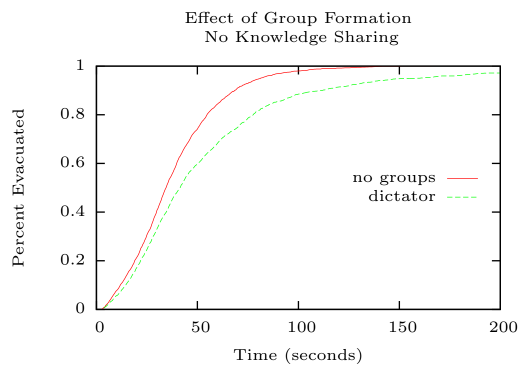

Figure 5.2: Comparison of static group formation on egress times without the sharing of knowledge using the following mechanisms: 1) No Groups and 2) Dictator

The last experiment for static groups compares the egress times of pedestrian egress when group formation occurs and knowledge sharing is allowed. The first route consensus mechanisms are tested pair-wise with the knowledge sharing mechanisms so we have the following combinations: 1) No Groups, 2) Dictator, 3) Least-cost, Private information (cost-prvt), 4) Least-cost, Public Information (cost-pblc, 5) Most-common, Private Information (coomon-prvt), and 6) Most-common, Public Information (common-pblc). As with the previous test, the pedestrians use the training route selection algorithm and have been trained 15 times in the building. This means that most agents know one or two exits and several ways to get there. As the cost of communication is likely to vary depending upon the circumstances, we consider two different communication costs, free (0 seconds per agent) and cheap (10 seconds per agent). As will be seen from the experiments, communication costs that are much greater than 10 seconds per agent are no longer effective, so we do not consider them.

Figure 5.3 compares the resulting egress times when knowledge sharing is free. When group members have their route-cost knowledge public, group formation leads to decreased egress times compared to no groups. Because all the route information is public between group members, the consensus mechanism has little effect on the egress time and only the cost-pblc mechanism is shown in the results. When group members keep their knowledge private, the consensus mechanism has a greater effect upon egress times. If the group uses the least-cost consensus mechanism, the egress performance is nearly as good as if the group had route-costs public among them and outperforms the egress time of individuals who do not form groups. If the group uses the most common consensus mechanism, egress performance is actually worse than not forming groups, as the common consensus mechanism is unable to capitalize on the information that may be had by only a minority of the group members. These results show that when knowledge is shared for free among group members, group formation is transformed from having a negative effect on egress times to having a positive effect on egress times.

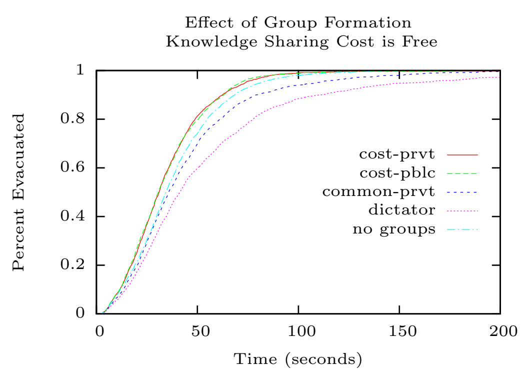

Figure 5.3: Comparison of static group formation on egress times using the following mechanisms: 1) No Groups, 2) Dictator, 3) Least-cost, Private information (cost-prvt), 4) Least-cost, Public Information (cost-pblc), and 5) Most-common, Private Information (common-prvt). Knowledge sharing is considered to be free.

Assuming knowledge sharing is free is, of course, an unrealistic assumption so Figure 5.4 compares the resulting egress times when knowledge sharing is cheap, incurring a cost of 10 seconds per group member. Under this assumption, having knowledge public among group members is no longer the most efficient solution. However, it can still improve the efficiency of egress as compared to no knowledge sharing. It improves the efficiency with which pedestrians exit as compared to the dictator mechanism after the 70% evacuated mark, and it improves efficiency for the last 6% of pedestrians as compared to the common-consensus mechanism without knowledge sharing. Using the least-cost route consensus mechanism with private knowledge (cost-prvt) is the most efficient technique. It does not incur the communication costs of publicly sharing all route knowledge but is still able to benefit from the individual knowledge of each pedestrian. In these experiments, cost-prvt is even more efficient than no groups. This is significant because it indicates that group formation does not always have to have a negative effect on egress times. If the knowledge can be used without explicitly being shared, the whole group can benefit.

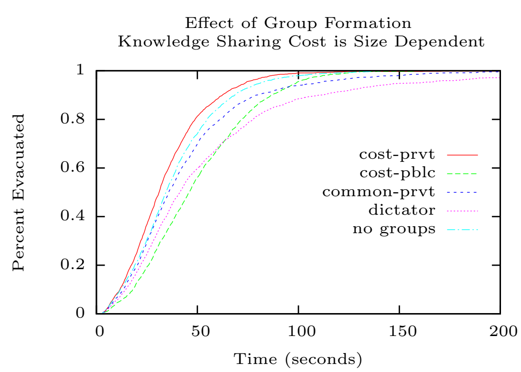

Figure 5.4: Comparison of static group formation on egress times using the following mechanisms: 1) No Groups, 2) Dictator, 3) Least-cost, Private information (cost-prvt), 4) Least-cost, Public Information (cost-pblc), and 5) Most-common, Private Information (common-prvt). Knowledge sharing is considered to cause a delay of 10 seconds per group member.

Besides the efficiency of each mechanism, we also consider several other statistics that indicate how well static groups maintain a close proximity (see Table 5.1). The two most relevant statistics deal with the spatial and temporal proximity maintained by groups while exiting the building. Spatial proximity is calculated as the average pair-wise distance be-

Table 5.1: Ranking of Route Consensus and Knowledge Sharing Mechanisms by Spatial Proximity

Rank	Mechanism	Temporal Proximity of Finish Point	Spatial Proximity of Finish Point
1	cost-pblc	7.078 s	2.33 m
2	common-pblc	7.52 s	2.74 m
3	cost-prvt	7.95 s	5.11 m
4	common-prvt	12.04 s	5.43 m
5	dicator	11.95 s	6.00 m
6	no group	26.82 s	26.63 m

tween the locations at which each group member finishes. Temporal proximity is defined as the amount of time elapsed between the successful egress of the first pedestrian of the group and the successful egress of the last pedestrian of the group. These two measures reflect how well group members maintain a close formation in both time and space. Table 3.1 shows the average spatial and temporal proximities for each combination of mechanisms considered. The results are ranked from closest spatial proximities to furthest spatial proximities. As would be expected, actively maintaining the group formation leads to closer proximities than no groups regardless of the mechanism used. Having knowledge information public among group members further increases the proximities group members are able to maintain because each member has the same knowledge, after the initial sharing has occurred. This might indicate one reason for sharing route information even with an increased communication cost.

5.4.2 Dynamic Groups

We perform similar experiments with dynamic groups to evaluate the effect of dynamic group formation on pedestrian egress time. The results obtained during these experiments indicate that, like static group formation, dynamic group formation tends to lead to slower egress times. The degree to which egress time is affected is dependent upon the number and size of the groups that form. When only a few small groups form, egress time is not significantly affected. However, as more groups are formed and as group size increases, the negative impact on egress times also increase. When the egress times for the stress and knowledge utility functions are compared, there does not seem to be a significant difference. It does not matter whether groups are formed because of insufficient knowledge or as a means to relieve stress regardless of the knowledge level; the biggest predictor of egress times is the number and size of groups formed.

5.5 Conclusions and Future Work

Pedestrian simulation is an important area of research with many applications. Until recently, group formation in pedestrian egress has largely been ignored. However, recent work has begun to address the issues that arise with group formation. In this paper, we have implemented a novel dynamic group formation technique that allows pedestrian groups to communicate, share knowledge, and reach a consensus regarding route selection. To our knowledge, this is the first such simulation model to address the issues of knowledge sharing and group consensus in pedestrian egress. We have shown that although recent literature emphasizes the negative impacts group formation can have upon egress times, positive incentives to group formation exist. Our simulation model predicts that sharing knowledge in pedestrian groups can help pedestrian maintain proximity with greater ease as well as improve egress times, compared to group formation without the sharing of knowledge. Additionally, group formation is found to be especially effective when groups do not explicitly share knowledge, but accept the least-cost route proposed by group members. Dynamic group formation has an impact on egress times similar to those found with static group formation. The number of groups formed and the size of the groups both have a greater impact on egress times than the reason for the group formation. Pedestrians can form groups to compensate for a lack of knowledge or as a means to reduce stress, but both reasons have similar impacts on total egress times.

As future work, we suggest that the communication costs for sharing route knowledge be evaluated empirically. We also suggest that similar experiments be performed with dynamic group formation to determine appropriate overhead costs for the group formation techniques.

CHAPTER 6

CONCLUSIONS AND FUTURE WORK

This work has made the following contributions to the field:

- The implementation of a new pedestrian simulation model PLEASE for modeling uncertainty in pedestrian knowledge, group formation, and information sharing.
- An evaluation of heuristic functions for predicting egress routes in a variety of real buildings.
- The novel use of reinforcement learning to simulate pedestrian knowledge.
- The use of reinforcement learning to utilize congestion information.
- The implementation of group consensus and knowledge sharing mechanisms to evaluate coalition formation costs and benefits.

The model proposed in this work represents an important step towards making pedestrian simulation models more accurately reflect the real world by simulating the uncertainty in pedestrian knowledge and allowing pedestrians in a group to collaboratively determine the egress route they should take. We have shown that assigning heuristic costs to decision points in the building can be an effective way to egress even when no knowledge of the building exists *a priori*. We have also shown that understanding congestion levels plays an important role in efficient pedestrian egress. In addition, we utilize reinforcement learning as an effective technique to provide pedestrians with unique, individualized knowledge about a building. Finally, group formation is found to have a significant impact on egress times, and interaction among group members can be both positive or negative depending upon the situation.

While answering several important questions in pedestrian egress simulation, this project leaves several potential directions for future work. Choosing an effective heuristic for evaluation of routes is highly dependent upon the building layout. Can we categorize buildings or, better yet, automatically learn to adjust the heuristic weights to effectively navigate unfamiliar building layouts? In addition, the communication model between pedestrians could be enriched so that additional information, such as the location of hazards, can be shared among pedestrians.

REFERENCES

[1] Antonini, G., Bierlaire, M., and Weber, M. Discrete choice models of pedestrian walking behavior. *Transportation Research Part B: Methodological 40*, 8 (2006), 667–687.

[2] Banerjee, B., Bennett, M., Johnson, M., and Ali, A. Congestion avoidance in multi-agent-based egress simulation. In *IC-AI* (2008), pp. 151–157.

[3] Barnes, M., Leather, H., and Arvind, D. Emergency evacuation using wireless sensor networks. *Local Computer Networks, Annual IEEE Conference on* (2007), 851–857.

[4] Büchner, S., Hölscher, C., and Strube, G. Path choice heuristics for navigation related to mental representations of a building. In *Proceedings of the European Cognitive Science Conference* (May 2007), Taylor & Francis, pp. 504–509.

[5] Busoniu, L., Babuska, R., and De Schutter, B. A comprehensive survey of multiagent reinforcement learning. *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on 38*, 2 (2008), 156–172.

[6] Claus, C., and Boutilier, C. The dynamics of reinforcement learning in cooperative multiagent systems. In *Proceedings of the National Conference on Artificial Intelligence* (1998), John Wiley & Sons LTD, pp. 746–752.

[7] Dalton, R. The secret is to follow your nose. *Environment and Behavior 35*, 1 (2003), 107–131.

[8] Dia, H. An agent-based approach to modelling driver route choice behaviour under the influence of real-time information. *Transportation Research Part C: Emerging Technologies 10*, 5-6 (2002), 331–349.

[9] Feuz, K. D., and Allan, V. Pedestrian route selection with imperfect knowledge. In *Proceedings of the 4th International Conference on Agents and Artificial Intelligence* (2012).

[10] Frankenstein, J., Büchner, S., Tenbrink, T., and Hölscher, C. Influence of geometry and objects on local route choices during wayfinding. In *Spatial Cognition VII*, C. Hölscher, T. Shipley, M. Olivetti Belardinelli, J. Bateman, and N. Newcombe, Eds., vol. 6222 of *Lecture Notes in Computer Science*. Springer Berlin / Heidelberg, 2010, pp. 41–53.

[11] Fridman, N., and Kaminka, G. Towards a cognitive model of crowd behavior based on social comparison theory. In *Proceedings of the National Conference on Artificial Intelligence* (2007), vol. 22, Menlo Park, CA; Cambridge, MA; London; AAAI Press; MIT Press; 1999, p. 731.

[12] Golledge, R. Path selection and route preference in human navigation: A progress report. *Spatial Information Theory A Theoretical Basis for GIS* (1995), 207–222.

[13] Goodrich, M., Stirling, W., and Frost, R. A theory of satisficing decisions and control. *Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on 28*, 6 (1998), 763–779.

[14] Gwynne, S., Galea, E., and Lawrence, P. The introduction of social adaptation within evacuation modelling. *Fire and Materials 30*, 4 (2006), 285–309.

[15] Gwynne, S., Galea, E. R., Lawrence, P. J., and Filippidis, L. Modelling occupant interaction with fire conditions using the buildingexodus evacuation model. *Fire Safety Journal 36*, 4 (2001), 327–357.

[16] Hassan, F., and Tucker, A. Using cellular automata pedestrian flow statistics with heuristic search to automatically design spatial layout. In *2010 22nd IEEE International Conference on Tools with Artificial Intelligence (ICTAI)* (Oct. 2010), vol. 2, pp. 32–37.

[17] Helbing, D. A mathematical model for the behavior of pedestrians. *Behavioral Science 36*, 4 (1991), 298–310.

[18] Helbing, D. A mathematical model for the behavior of individuals in a social field. *Journal of Mathematical Sociology 19*, 3 (1994), 189–219.

[19] Helbing, D., and Johansson, A. Pedestrian, crowd and evacuation dynamics. In *Encyclopedia of Complexity and Systems Science*. Springer, 2009, pp. 6476–6495.

[20] Helbing, D., Schönhof, M., Stark, H., and Holyst, J. How individuals learn to take turns: Emergence of alternating cooperation in a congestion game and the prisoner's dilemma. *Advances in Complex Systems 8*, 1 (2005), 87–116.

[21] Hillier, B. *Space is the Machine.* Cambridge University Press, 1996.

[22] Hochmair, H., and Frank, A. Influence of estimation errors on wayfinding-decisions in unknown street networks–analyzing the least-angle strategy. *Spatial Cognition and Computation 2*, 4 (2000), 283–313.

[23] Hoogendoorn, S., and Bovy, P. Pedestrian route-choice and activity scheduling theory and models. *Transportation Research Part B: Methodological 38*, 2 (2004), 169–190.

[24] Hoogendoorn, S. P., and Daamen, W. Pedestrian behavior at bottlenecks. *Transportation Science* (2005), 147–159.

[25] Integrated Environmental Solutions Limited. Simulex - simulation of occupant evacuation. http://simcosm-india.com/pdf/Simulex.PDF, 2008. Retrieved July 28, 2011.

[26] Kaelbling, L., Littman, M., and Moore, A. Reinforcement learning: A survey. *Journal of Artificial Intelligence Research 4* (1996), 237–285.

[27] Koh, W. L., Lin, L., and Zhou, S. Modelling and simulation of pedestrian behaviours. In *PADS '08: Proceedings of the 22nd Workshop on Principles of Advanced and Distributed Simulation* (Washington, DC, USA, 2008), IEEE Computer Society, pp. 43–50.

[28] Kray, C., Kortuem, G., and Krüger, A. Adaptive navigation support with public displays. In *Proceedings of the 10th international conference on Intelligent user interfaces* (New York, NY, USA, 2005), IUI '05, ACM, pp. 326–328.

[29] Kretz, T., Hengst, S., and Vortisch, P. Pedestrian flow at bottlenecks - validation and calibration of vissim's social force model of pedestrian traffic and its empirical foundations. Tech. Rep. arXiv:0805.1788, May 2008. Comments: Accepted for presentation at International Symposium of Transport Simulation 2008 (ISTS08) in Gold Coast, Australia.

[30] Lewin, K. *Field theory in social science: selected theoretical papers.* Harper & Row, New York, 1951.

[31] Littman, M. Markov games as a framework for multi-agent reinforcement learning. In *Proceedings of the eleventh international conference on machine learning* (1994), vol. 157, Citeseer, p. 163.

[32] Littman, M. Friend-or-foe Q-learning in general-sum games. In *Machine Learning-International Workshop Then Conference* (2001), pp. 322–328.

[33] Luber, M., Stork, J., Tipaldi, G., and Arras, K. People tracking with human motion predictions from social forces. In *2010 IEEE International Conference on Robotics and Automation (ICRA)* (2010), IEEE, pp. 464–469.

[34] Luke, S., Cioffi-Revilla, C., Panait, L., Sullivan, K., and Balan, G. Mason: A multiagent simulation environment. *Simulation: Transactions of the Society for Modeling and Simulation International 81*, 7 (2005), 517.

[35] Moussaïd, M., Perozo, N., Garnier, S., Helbing, D., and Theraulaz, G. The walking behaviour of pedestrian social groups and its impact on crowd dynamics. *PLoS One 5*, 4 (2010), e10047.

[36] Müller, J., Jentsch, M., Kray, C., and Krüger, A. Exploring factors that influence the combined use of mobile devices and public displays for pedestrian navigation. In

Proceedings of the 5th Nordic Conference on Human-computer Interaction: Building Bridges (New York, USA, 2008), NordiCHI '08, ACM, pp. 308–317.

[37] Murakami, Y., Minami, K., Kawasoe, T., and Ishida, T. Multi-agent simulation for crisis management. In *2002. Proceedings. IEEE Workshop on Knowledge Media Networking* (2002), IEEE, pp. 135–139.

[38] Ozel, F. Time pressure and stress as a factor during emergency egress. *Safety Science 38*, 2 (2001), 95–107.

[39] Pan, X. *Computational Modeling of Human and Social Behaviors for Emergency Egress Analysis*. PhD thesis, Stanford University, Stanford, California, 2006.

[40] Pan, X., Han, C. S., Dauber, K., and Law, K. H. A multi-agent based framework for the simulation of human and social behaviors during emergency evacuations. *AI & Society 22*, 2 (2007), 113–132.

[41] PTV AG. http://www.ptvag.com/software/transportation-planning-traffic-engineering/software-system-solutions/vissim-pedestrians/, 2011. Retrieved April 4, 2011.

[42] Qingge, J., and Can, G. Simulating crowd evacuation with a leader-follower model. *IJCSES 1*, 4 (2007), 250.

[43] Qiu, F. *A Framework for Group Modeling in Agent-Based Pedestrian Crowd Simulations*. PhD thesis, Georgia State University, Atlanta, Georgia, 2010.

[44] Qiu, F., and Hu, X. Modeling group structures in pedestrian crowd simulation. *Simulation Modelling Practice and Theory 18*, 2 (2010), 190–205.

[45] Roughgarden, T. The price of anarchy is independent of the network topology. In *Proceedings of the Thiry-Fourth Annual ACM Symposium on Theory of Computing* (2002), ACM, pp. 428–437.

[46] Roughgarden, T., and Tardos, É. How bad is selfish routing? *Journal of the ACM (JACM) 49*, 2 (2002), 236–259.

[47] Russell, S., and Norvig, P. *Artificial Intelligence: A Modern Approach*. Prentice Hall, 2010.

[48] Santos, G., and Aguirre, B. E. A critical review of emergency evacuation simulation models. In *Building Occupant Movement during Fire Emergencies* (Gaithersburg, Maryland, June 2004).

[49] Shehory, O., and Kraus, S. Methods for task allocation via agent coalition formation. *Artificial Intelligence 101*, 1-2 (1998), 165 – 200.

[50] Soh, L., and Tsatsoulis, C. Satisficing coalition formation among agents. In *Proceedings of the First International Joint Conference on Autonomous Agents and Multiagent Systems: Part 3* (2002), ACM, pp. 1062–1063.

[51] Thompson, P., and Marchant, E. Testing and application of the computer model simulex. *Fire Safety Journal 24*, 2 (1995), 149–166.

[52] Wahle, J., Bazzan, A., Klügl, F., and Schreckenberg, M. Decision dynamics in a traffic scenario. *Physica A: Statistical Mechanics and its Applications 287*, 3-4 (2000), 669–681.

[53] Wahle, J., Bazzan, A., Klügl, F., and Schreckenberg, M. The impact of real-time information in a two-route scenario using agent-based simulation. *Transportation Research Part C: Emerging Technologies 10*, 5-6 (2002), 399–417.

[54] Wierzbicki, A. A mathematical basis for satisficing decision making. *Mathematical Modelling 3*, 5 (1982), 391–405.

[55] Yang, L., Zhao, D., Li, J., and Fang, T. Simulation of the kin behavior in building occupant evacuation based on cellular automaton. *Building and Environment 40*, 3 (2005), 411–415.

[56] Zhao, D., Yang, L., and Li, J. Occupants' behavior of going with the crowd based on cellular automata occupant evacuation model. *Physica A: Statistical Mechanics and its Applications 387*, 14 (2008), 3708–3718.

[57] Zheng, X., Zhong, T., and Liu, M. Modeling crowd evacuation of a building based on seven methodological approaches. *Building and Environment 44*, 3 (2009), 437–445.