INCORPORATING GROUPS, COLLECTIVE BEHAVIOR, AND INFORMATION VISUALIZATION IN AGENT-BASED MODELS OF EVACUATION

by

Eric Best

A dissertation submitted to the Faculty of the University of Delaware in partial fulfillment of the requirements for the degree of Doctor of Philosophy in Disaster Science and Management

Spring 2013

INCORPORATING GROUPS, COLLECTIVE BEHAVIOR, AND INFORMATION VISUALIZATION IN AGENT-BASED MODELS OF EVACUATION

by

Eric Best

Approved:

Leland Ware, J.D.
Director of the School of Public Policy & Administration

Approved:

George H. Watson, Ph.D.
Dean of the College of Arts & Sciences

Approved:

James G. Richards, Ph.D.
Vice Provost for Graduate and Professional Education

I certify that I have read this dissertation and that in my opinion it meets the academic and professional standard required by the University as a dissertation for the degree of Doctor of Philosophy.

Signed:

____

Benigno Aguirre, Ph.D.
Professor in charge of dissertation

I certify that I have read this dissertation and that in my opinion it meets the academic and professional standard required by the University as a dissertation for the degree of Doctor of Philosophy.

Signed:

____

James Corbett, Ph.D.
Member of dissertation committee

I certify that I have read this dissertation and that in my opinion it meets the academic and professional standard required by the University as a dissertation for the degree of Doctor of Philosophy.

Signed:

____

Sue McNeil, Ph.D.
Member of dissertation committee

I certify that I have read this dissertation and that in my opinion it meets the academic and professional standard required by the University as a dissertation for the degree of Doctor of Philosophy.

Signed:

____

Joseph Trainor, Ph.D.
Member of dissertation committee

ACKNOWLEDGMENTS

I would like to thank my advisor, Dr. Benigno Aguirre for his guidance, support, and mentorship. I would also like to acknowledge the esteemed members of my Ph.D. committee for their assistance, direction, and critical input: Dr. James Corbett, Dr. Sue McNeil, and Dr. Joseph Trainor.

I would like to recognize the National Science Foundation, the source of funding support for the project underlying in this dissertation.

My research was conducted with the support of a fantastic group of undergraduate research assistants who made significant contributions to the NSF project; Etornam Banini, Christina Dalton, Michael Lopez, Adam Orescan, and Shawn Reynolds.

I would like to thank all of the model builders I interacted with over the last four years. While social behavior has yet to play an important role in the building evacuation modeling community, I am hopeful that this situation will change. Everyone I communicated with expressed a tireless desire to make the best models possible with the information, funds, and computing power available today.

Finally, I would like to thank my family for encouraging me to return to graduate school when it did not seem to make sense, and for their constant love and support

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

ABSTRACT

This dissertation is intended to advance research in building evacuation modeling through the introduction of detailed social groups, collective behavior, and improvements in information visualization. The model built as a part of this dissertation makes significant original contributions to both input and output of building evacuation models.

Regarding inputs, this work prototypes new ways to catalog social groups, leadership, and the concept of supra force – a combination of high density, contraflows of crowds, and environment. The central difference between this model and previous efforts is the role of group affiliations. The effort resulting from this dissertation, SocEvac, creates a three-layer decision tree for most agents, who have to balance individual and group responsibilities while attempting to avoid supra force. This interaction of individually optimal exit paths and social responsibilities creates significantly more contraflow situations as agents attempt to locate and evacuate with their loved ones. These contraflows impede efficient evacuation, helping to explain scenarios such as the Station nightclub evacuation; where there were significantly more fatalities then would have been expected based on population and number of exits.

On the output side, this work creates new methods to examine simulation models in real-time, and suggests new methods of measurement to determine what makes an accurate model. The real-time visualization methods allow for researchers to quickly understand what is happening while a model is running. These visualizations allow for users of a simulation to control what features they want to highlight in a model in real-time. The new methods of output measurement center around tracking agents as

individuals, cataloging outcomes of agents each modeling a real-world counterpart complete with demographics and relationships. By transitioning away from aggregate population tracking and focusing on individuals, it is now possible to compare models to an entire evacuation narrative instead of only attempting to recreate end results.

These improved inputs and outputs result in the most descriptive model to date of population movements during the Station nightclub fire in 2003. More than ten years after the fire, I believe the SocEvac model can finally begin to explain the complex events that led to the high fatalities and unconventional exit paths of the evacuation.

While this dissertation focuses on one scenario, the underlying program can be used to model almost any building evacuation. This platform is designed to inspire other model builders to consider adding group social behaviors to models.

PREFACE

As a researcher interested in merging separate areas of work in social science evacuation research and building evacuation modeling, the interdisciplinary disaster science and management program at the University of Delaware was the perfect outlet for this undertaking. Although this is a dissertation that explains the development of a model, it is really a commentary about the state of building evacuation modeling. While many fields have made huge advances in building evacuation modeling and human behavior analysis, these fields rarely overlap.

While the situation regarding evacuation models and social behavior looks more promising today than it did when this project began, there is still an almost complete separation between those that study human interaction and behavior and those that model evacuations or egress. The model outlined in this dissertation is meant to serve as a prototype to inspire other modeling efforts incorporating behavioral advances in social science. This model is intended to be the beginning on an enhanced modeling effort, not the end of implementation of behavior in modeling. My hope is that model builders will consider this and other efforts to redefine what “social behavior” means in evacuation models.

Chapter 1

STATEMENT OF PROBLEM: IMPROVING EVACUATION MODELING BY INCORPORATING SOCIAL RELATIONSHIP BEHAVIOR

This dissertation is an attempt to bridge the gap between social science evacuation research and building evacuation modeling. Computational modeling of evacuation involves creating every element of virtual worlds. Many models are built from scratch, giving modelers complete control over the environments they create. This freedom allows evacuation model builders to create complex environments with detailed building layouts, hazards, weather, and seemingly countless other parameters. However, despite an ability to be masters of their virtual domains, evacuation model builders rarely account for detailed human behavior and interaction. This is a significant oversight in a field that prides itself on attention to detail. Despite extensive social scientific research about human social behavior, most evacuation models do not incorporate even rudimentary social elements. Most models designed to replicate mass evacuations (where many people find themselves evacuating with friends or loved ones) completely disregard these social bonds.

I believe that model builders and social scientists working together can improve multiple disciplines, including disaster management, by building models that account for established social science theories. Merging established social science theory with widely accepted simulation coding methods has the potential to create more meaningful and interpretable results from models, as well as to contribute to the growing field of hazard modeling.

1.1 Applying Theory to Modeling

This dissertation is an attempt to build upon a previous body of work related to modeling building evacuations by incorporating some elements of social science theory. In 2009, the Disaster Research Center (DRC) developed a conceptual map of what a properly specified model should look like. This model included ecology, hazards, heterogeneous individual profiles of victims, group social relationships, and the effects of crowd densities in confined spaces.

Current building evacuation models did not focus on social relationships, and did not even include the ability to account for them. The first-generation model developed in this collaboration was a true interdisciplinary project, combining research from the multiple areas of expertise. The results of the model suggest that the initial assumptions could be used to create more accurate building evacuation models.

After the first-generation model was completed, it was clear that a better, second-generation model was needed. The model would be generalizable, and be accessible and understandable to users from multiple disciplines. These changes and many others were combined to create a second-generation building evacuation model accounting for social behavior, the objective of this dissertation.

The development of these models and the analysis of observed patterns can be judged as an original contribution to disaster science and management. The results of the models accounting for social behavior raise some interesting questions, such as the importance of distance and clearance times in building evacuation models versus what really happens when a group of people attempts to evacuate a building. While this research is intended to complement and influence current building evacuation models instead of replace them, it will show the importance of social science concepts as tools to improve evacuation models.

In a way, this dissertation is an answer to the challenge posed in Manuel Torres's dissertation, where he suggests that conventional model design falls short of returning results that make sense when looking at victims of the Station nightclub fire as a network of related individuals instead of a homogenous population (Torres 2010).

1.1.1 The Current State of Building Evacuation Modeling

In certain areas of both science and practice, it is important to examine real-life situations in simulated environments. There are occasions where injury, disease, emotional or physical harm, or other negative impacts that generally cannot be explored in a live research setting due to cost, legal implications, or professional ethics guiding human or animal research, are nevertheless of substantive significance for human safety. People interested in these occurrences in natural settings generally can only look at them after the fact, which completely prevents the ability to ask "what if..." questions that can be addressed through experimentation in many other fields. Researchers and practitioners are unable to run experiments in these other settings, so they rely on computer simulations to validate results or attempt to predict future behaviors (Banks et al. 2004). Because much of disaster science focuses on extreme scenarios, modeling and simulation will be an increasingly important part of the field, offering a bridge between theory and practice.

The topics covered in disaster science frequently cannot be tested in real-world experiments. It is not possible to evacuate a city just to study the results, or intentionally infect members of an experimental group to measure an increase in fatality rate. Topics that cannot be studied experimentally in the real world often involve important issues that need to be addressed, making simulation experiments necessary. But if the results of simulation are going to be used to create or alter policy pertaining to major societal issues, the simulation models' validity becomes very important. Simulating humans is very difficult, and sacrifices to the quality and robustness of models are made to allow quick human simulation (Pidd 2010). Unfortunately, what is included in a model can be even more harmful than what is left out. For example, despite decades-long research about the rarity of "panic" in evacuation scenarios, the concept is still included as the main engine in some models of evacuation (Abdelhak, Daude, and Olivier 2009, Barnshaw, Letukas, and Quarantelli 2008). Panic and antisocial behavior are two of E.L. Quarantelli's top six "disaster myths" (Quarantelli 2008), yet individual behavior and panic scenarios are often found in building evacuation models, where it is posited, against the results of research in collective behavior, that in leader/follower situations agents do not actually care about each other.

At this time, building evacuation modeling is dominated by the disciplines of civil engineering and computer science. While these disciplines develop state-of-the-art build evacuation models, they do not leave a great deal of room for social scientists to make contributions. Many evacuation models concentrate on optimization of exit strategies, which is good in theory, but relying too much on optimization often prevents multiple solutions from being considered. Too often in modeling, the "best" solution is the only one considered at the expense of all others. At one point during my dissertation writing, the editor of a journal asked me to review a paper that introduced a building evacuation algorithm that concluded that multiple exit signs should be removed from buildings, leaving only one clearly signed evacuation path from any building point. While the author certainly did not mean to do harm to building occupants, there are very obvious flaws in only considering the "best" choice in high-stakes events like building evacuations. Very often the theoretical solution that is the best option is not the best practical solution.

This work creates models that offer multiple possibilities, multiple solutions, and multiple outcomes for even simple scenarios. In a sense, it replicates a small percentage of the chaos that exists in the real world. This requires reassessing or abandoning many of the assumptions that are currently in "state-of-the-art" evacuation modeling. Models cannot rely solely on which exit is closest or which time is fastest to have avatars choose building exits. While these are widely accepted conventions, they do not make it possible to create models that replicate real-world results where evacuees delay exit, take "inefficient" paths, or worry about things other than their own safety and speedy evacuation, such as spouses or children. Nevertheless, adding some of these social factors will likely make it possible to better replicate real-world evacuation results.

Researchers create models in an attempt to understand previous occurrences, predict future performance, and answer "what if..." questions. Usually, when a disaster occurs there are multitudes of failures, and there is always a desire to know whether things would have worked out for the better if there had been situational differences. Most of the time, disaster models are ultimately concerned with the safety of people, so computer simulations attempt to approximate physical environments and human outcomes. As mentioned by other researchers exploring the Station scenario, conventional modeling platforms are better at prediction than replication, because attempts to make the models as efficient as possible makes them bad at highlighting the inefficiencies in real-world evacuations that almost always occur (Spearpoint 2012). The model in this dissertation was designed to account for inefficiency, making it fundamentally different from many commercial efforts.

In some instances, models of disaster scenarios suffer from correlation versus causation issues; models only track results, not what happens during the simulations. In overly rigid models, inputs predetermine the results. For example, a trend in evacuation models is to use previous models as benchmarks of accuracy instead of real-world results, which is not an ideal method for correct simulations (Aguirre et al. 2011a). Ideally, model builders should create models of past events based on detailed data about the event, which offer much more precise standards for judging outcomes than the output of previous models. Nevertheless, other problems are created that demand attention; too much detailed input data predetermine results, so that models cease to simulate and simply confirm input bias. Relying on past simulations instead of actual events can lead to amplifications of inaccuracies between model iterations, and relying on the wrong real-world data can lead to models that predispose the outcome.

There has been a growing recognition amongst computer modelers and end users of the need to incorporate social behavior to improve model accuracy (Cherif and Djedi 2005, Helbing, Johansson, and Al-Abideen 2007, Moussaid et al. 2010, Galea 2004). Historically, most of the building evacuation modeling platforms developed by the civil engineering and computer science disciplines focus on one of two calculation methods for exit paths for agents: closest distance and most efficient (or fastest) clearance. These two methods have served model builders well for decades, but there are significant shortcomings with both methods.

Closest exit models, a category containing a variety of flow models, create conditions that make it difficult to gain useful knowledge about human environments. If all agents in a model simply select the exit that is closest to them, a model becomes a simple set of physics equations, dictating how many agents can pass through a choke point such as a door per second. When all agents are taking the closest path out of an environment, it is impossible to have contraflows (Ehibara, Ohtsuki, and Iwaki 1992, Fang et al. 2010).

Efficient clearance calculations are an improvement, but still leave much to be desired when attempting to model human behavior. One prominent example of a fastest clearance model uses the A* algorithm, first developed in the 1960s, which calculates the most efficient path for each exit in a building evacuation model (Hart, Nilsson, and Raphael 1968). Fastest clearance calculations technically can allow for contraflows, although these situations very rarely occur. Some platforms create hybrids of these two exit designs (Tissera, Printista, and Luque 2012).

More recently, model builders attempt other types of evacuation models, such as efforts based off of interview data about how people evacuate (Li, Tang, and Simpson 2004), or cellular automata (Pelechano and Malkawi 2008). Often, new platforms claim to model social behavior, but a more thorough analysis finds that there is little or no diversion from closest exit or fastest clearance calculations (Fang et al. 2010, Tissera, Printista, and Luque 2012). If social behavior is accounted for, oftentimes models rely on concepts not supported in literature, such as widespread panic (Ren, Yang, and Jin 2009), groups that move as one unit (Johnson and Feinberg 1997), and competitive social behavior (Pan et al. 2007). To date these efforts are rarely contextualized within greater social science theory, and there is acknowledgement that further efforts are required to represent human social behavior, although until now this is done within disciplines instead of between them (Kuligowski and Gwynne 2010, Pelechano and Malkawi 2008). This analysis was conducted using all available documentation, but it should be noted that some of the described platforms may possess more or less functionality than depicted if not specified in available literature (Kukla et al. 2001, Tsai et al. 2011, Mott MacDonald 2012, Spearpoint 2012, Gwynne et al. 2000, Manocha and Lin 2012, Pan et al. 2006, ARA 2012, Kuligowski and Peacock 2005, Kuligowski and Gwynne 2010, Haron et al. 2012). Below is a brief analysis of a selection of these platforms, chosen based on the capabilities they possess.

Evacuation Simulation with Children, Authorities, Parents, Emotions, and Social Comparison (ESCAPES) was created by the TEAMCORE lab at the University of Southern California. ESCAPES is structured in a four level action set: different agent types, emotional interactions, informational interactions, and behavioral interactions (Tsai et al 2011). This model incorporates basic social behavior as well as decision theory, however this is a strictly disciplinary effort involving a team made up of only computer scientists. There are a number of differences in the modeling approach used by the ESCAPES team and the approach our team is proposing, including the quantification of group dynamics and leadership, as well as the importance of changing environments (hazards and agents) in simulations. The social behaviors and interdependencies between individuals, groups, and a community during crisis situations are extremely complex problems. ESCAPES, despite its capabilities, still has limitations related to group understanding and interactions.

buildingEXODUS is an ABM created by the Fire Safety Engineering Group at the University of Greenwich (Gwynne et al. 2000, 2006). buildingEXODUS is a commercially available evacuation simulation model that is widely utilized in the United Kingdom. It consists of six sub models that interact with one another during an evacuation simulation: Occupant, Movement, Behavior, Toxicity, Hazard and Geometry sub models. Building space is modeled as a two dimensional spatial grid made of up of nodes and arcs along which agents travel between adjacent nodes. The motion of agents and their physical interaction with one another is therefore not as natural as in MASSEgress. The model takes into account various characteristics pertaining to each agent, including physical, psychological and positional values. buildingEXODUS also takes into account fire/agent interaction using computational fluid dynamic models that track the various elements of fire hazard. As with the other platforms, buildingEXODUS does not take into account the effects on evacuation of extreme crowd densities, and its treatment of group effects is limited (Gwynne et al. 2006).

MASSEgress is an ABM developed by Pan et al. (2006) for modeling emergency evacuation. Each agent in a MASSEgress simulation observes the surrounding environment and makes independent decisions based on these observations. Agents employ instinct rules, past experiences, social rules and rational inference to choose a behavior type (Pan, 2006). This complexity allows agents to independently choose an escape route at the micro level. It results in complex behavior like cooperative queuing and competitive herding near exits. MASSEgress does not model extreme crowding when agent responses occur in high-density situations. In addition, while it does account for some limited group effects, such as negotiation within a group to determine priority at an exit point, MASSEgress does not explicitly account for group effects that reflect the varying strengths of social ties among agents.

Researchers in the Department of Computer Science at the University of North Carolina Chapel Hill (UNC) created an NSF-funded project, "Interactive Large-Scale Crowd Simulation" (Manocha and Lin 2012). This is a platform that allows for the simulation of tens of thousands of agents in and out of buildings. This platform allows for heterogeneous crowds as well as agents, and simulates emergent and group behaviors. While this is another platform attempting to simulate large-scale gatherings of people, it does not appear to account for our definition of social behavior, the importance of relationships and unique knowledge. The social behaviors and interdependencies between individuals, groups, and a community during crisis situations are an extremely complex problem. The UNC development team acknowledges the need to consider crowd validation and analysis (Lin et al 2010).

Building evacuation models implementing social behavior are not limited to academic efforts. Applied Research Associates (ARA) offers a software platform called E-Sim (Event Simulator) to model building evacuations. This platform is capable of modeling large-scale buildings. E-Sim includes heterogeneous agents with "Agent behavior based on optimal choice or distribution of statistical human behavior" (ARA 2012). While ARA is a company with a great deal of experience in environmental simulation, it does not highlight the importance of social behavior in evacuations.

This dissertation will present a modeling format and analysis suite developed with these platforms in mind, using validity specifications created from sufficient event data, rather than from benchmarking older models or inserting probabilities. The data that are used are limited to known quantities at the inception of a hazard, not information only known after the fact (such as which exit an individual used.

Despite these more recent social behavior efforts, much of hazard modeling focuses on physics. In contrast, this work focuses on modeling behavior in evacuations using social characteristics, which as of now is an underserved aspect of modeling. Hazard modeling is a mature field, with a variety of platforms devoted to damage modeling, hazard spread, and hazard likelihood, among others (Cardona et al. 2008, Park, Ang, and Wen 1985, Pinelli et al. 2004, Stirling, McVerry, and Berryman 2002). Government entities like the National Institute of Standards and Technology employ model builders to create complex and thorough hazard models, including a model of the fire spread in the Station (Grosshandler et al. 2005). Engineers and architects frequently require extremely accurate building design systems, and there are numerous physics engines that can replicate buildings. For building dimensions, computer aided design (CAD) systems provide a standardized way to design computer representation of buildings that are accurate to scale (CADDPrimer 2011). Nevertheless, because the focus of this dissertation is on avatar behavior, and because the Station did not collapse during the evacuation period, there is no physics engine to predict structural changes in the sample building in these models. Since changes in the building pale in comparison to the changes produced by the smoke and fire (the hazards at the Station), the model will use a simple CAD system that inputs the dimensions of the building only.

This dissertation presents a computer platform accounting for individual behavior, social relationships or group behavior, the spread and effects of hazards, and the specific characteristics of the space the evacuees had to negotiate as they exited the building. It is a platform designed to allow researchers to isolate the effects of different types of human behavior. It does not attempt to replicate the findings in evacuation modeling literature that suggests that door size is the most important exit attribute, but instead pursues the idea that social relationships and small groups (identified by real-life evacuees) are key factors in evacuation decisions (Ha and Lykotrafitis 2012, Wang et al. 2011). Thus, the models presented are examples of interdisciplinary science using behavioral traits from the social sciences to guide creation of computer code.

A strong understanding of computer science is required to build simulations, a solid understanding of social science is helpful to create systems that accurately display human behavior. This dissertation attempts to create evacuation models that account for hazards, environmental factors, individual traits, group interaction, and the forces of crowds.

However, it quickly became obvious that creating behavior models is hardly limited to the knowledge obtained in a single discipline. For now, understanding of programming concepts is obviously required, but that could change in the future with increasingly friendly programming environments (Sengupta 2011), including the effort resulting from this dissertation. In order to increase the accessibility of building evacuation models, the model in this dissertation can be created and commanded by users that do not know how to code at a high level.

For academics interested in building evacuation models, there are many commercial software packages designed to assist with quick model creation, but they are often limited by their intent or design (Castiglione 2010). For instance, a commercial model of fire propagation usually would not include the possibility for other hazards to occur at the same time. Many commercial software packages, such as the Vissim platform, require users to purchase components piecemeal, so model builders are limited by the options their firms purchase (PTV 2012). In many cases, both the programs themselves and their output are considered proprietary, so it is difficult to be able to peer-review the modeling components that are used to make policy. This dissertation hypothesizes that group behavior elements, commonly referenced in interviews with survivors of disasters, are an important part of any successful evacuation model, and provides a software framework for others to use these code elements (Aguirre, El-Tawil, et al. 2011, Quarantelli 1960, Weller and Quarantelli 1973).

1.2 Incorporating Social Science in Models of Evacuation

A few model builders have already attempted to incorporate social behavior elements or at least identified the need to do so (Gwynne, Galea, and Lawrence 2006, Pan et al. 2006). Despite the inclusion of some social behavior, these models rely on questionable notions such as crowd panic (Helbing, Farkas, and Vicsek 2000). In contrast, my intention is to create a class of models that include social behavioral elements using accepted social science theories that account for social relationships and density in crowds. Following earlier research, an extensive review of modeling platforms by Aguirre et al. 2011a found that true social behavior and relationship concepts (such as caring about other people, represented in the model as avatars) are not used in agent-based models of evacuation (Kuligowski and Peacock 2005, Santos and Aguirre 2004). While model builders cannot implement all desired traits, it is not beneficial to default to an approach that does not consider behavioral elements, because these results will be inaccurate (Pan 2010). As illustrated below, adding even simple group dynamics make navigation through rendered physical environments significantly more difficult.

1.2.1 Issues With Group Behavior in Physical Environments

Before any coding begins, there are conceptual issues to address regarding modeling of groups navigating physical environments. Some of my original conceptual drawings are presented in Figures 1 to 5 to highlight potential issues addressed in the models.

Most of the work after basic coding involved creation of rules for group interaction, and modification of these rules after observing the behavior of avatars. This section highlights some of the problems model builders must address when accounting for group interactions.

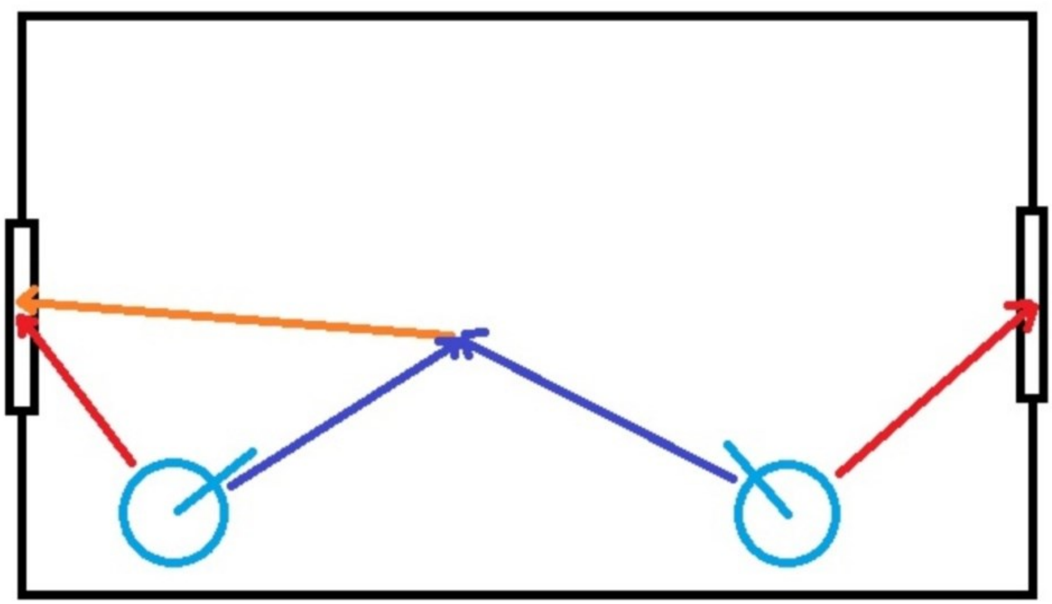

Figure 1.1: Two Agent Exit Selection

Figure 1.1 is a basic diagram of a conflict in exit options between two agents that are in a group. Agents in a group would prefer to exit together (the orange line), but their individual optimal exit paths are different (the red lines). Two agents acting in concert is one of the simplest group behavior scenarios, yet this example already shows some of the inherent difficulties and judgments required to have group behavior in a simulated physical environment (Fang et al. 2010). If the group members prioritized meeting each other over exiting, their paths would be similar to the purple lines. In my current model, the group leader would choose an exit independent of the other agents, and the followers would then attempt the same route trying to move with their group.

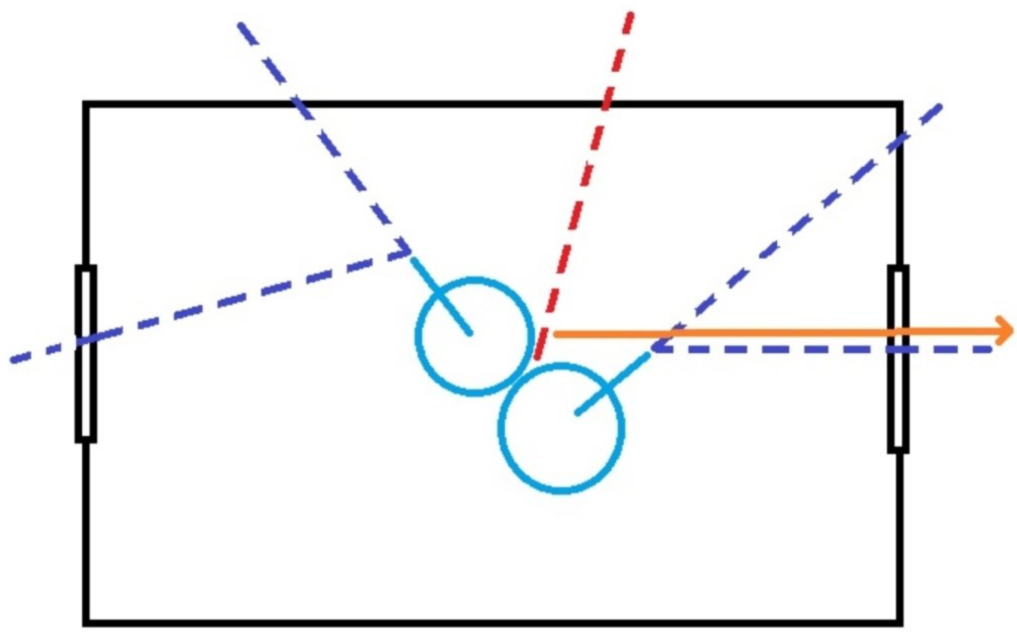

Figure 1.2: Exit Angle Calculation

Figure 1.2 shows a pair of agents and their minimum angle calculations to their individually optimal exits. Regardless of the behavior that is used, each avatar is programmed to make these calculations since lateral or angular movement is required before each avatar can use forward movement to exit. The purple dashed lines represent minimum angles to exits for each avatar. The red dashed line represents a bisection of the angle between the directions that the two agents are facing, essentially the fastest way that both agents can move together in the same direction. My hypothesis of how a distributed group would calculate an exit point would be to examine that bisected angle and its angular and linear distance from each exit, and then rank each choice. The rightmost exit would clearly be ideal, since less angular movement would be required, and the exit does not appear to be any farther than the exit on the left. However, if the leftmost agent is the "leader" in my current model, the group would take the exit on the left side.

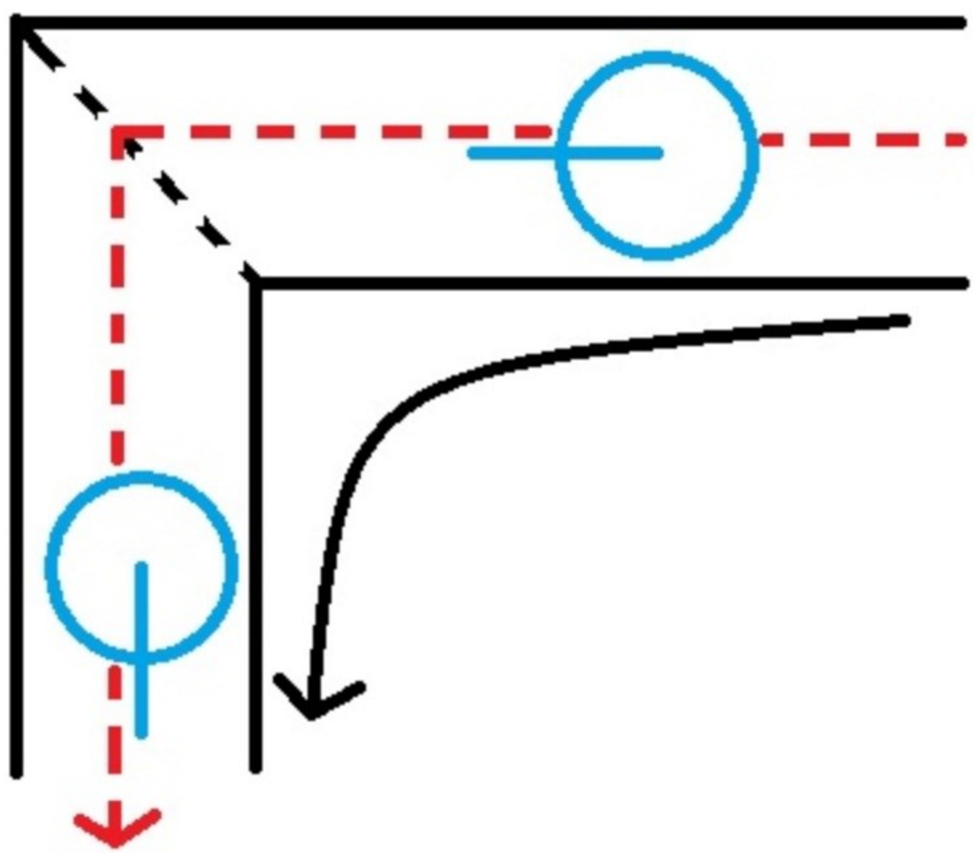

Figure 1.3: Line of Sight

Figure 1.3 illustrates a situation where agents moving in a group would have to intuitively understand that other agents in their group cannot see the exit. At this time, followers may abandon a group leader if they are unable to see them for a defined period.

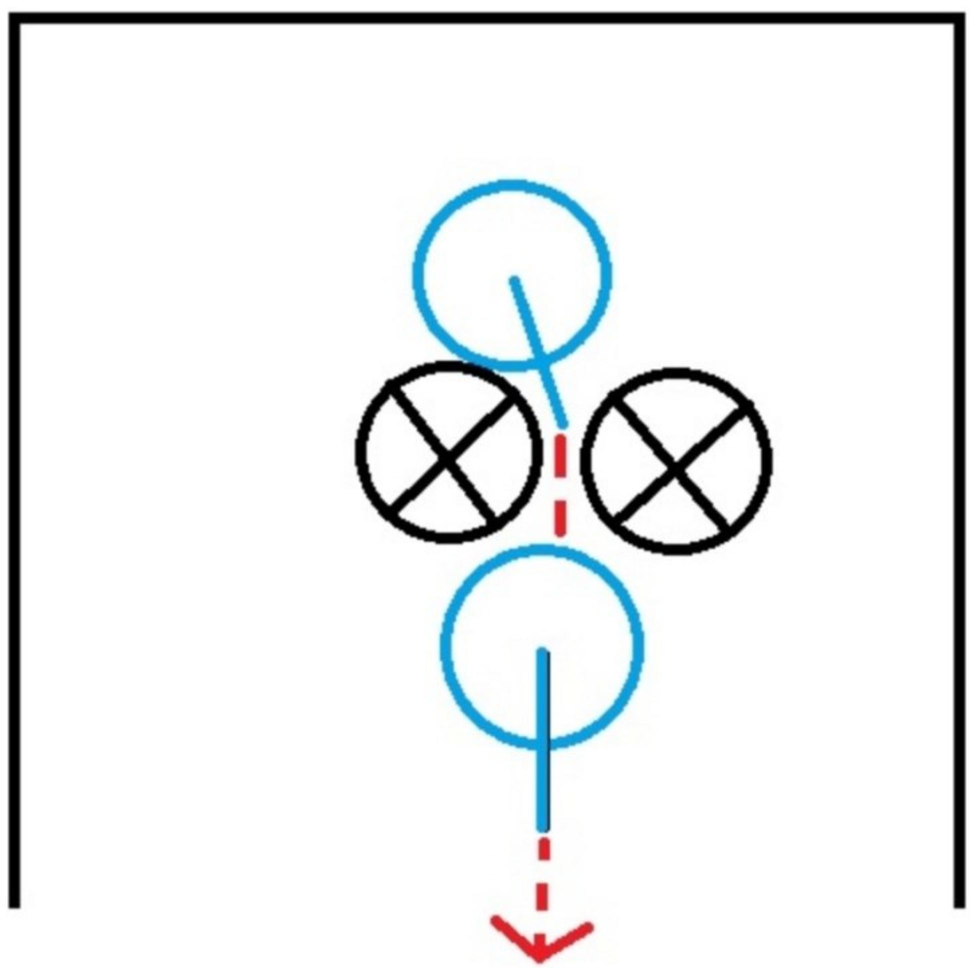

Figure 1.4: Overcoming Obstruction

Figure 1.4 illustrates a scenario that comes up often, when one or more members of a group have a desired exit path blocked by an obstruction (usually other avatars, especially inactive ones represented here by crossed circles). Eventually, it would be desirable to have groups where agents ahead of the obstruction would have the option to wait or to attempt to remove the obstruction in addition to the option to keep moving without relying on a time or distance equation. Weightings for each behavior choice would vary based on the type of group (intimate, social, etc) and the environmental conditions.

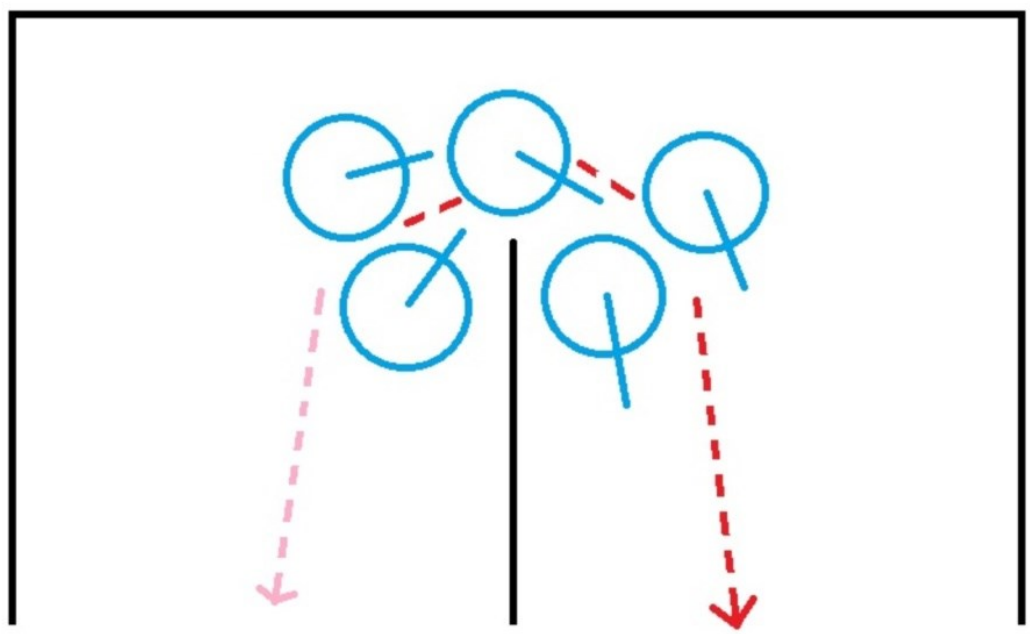

Figure 1.5: Inconsistent Goals

Figure 1.5 illustrates a group moving around a barrier, where there are exits on either side. This presents an occasion where there is clearly an inconsistent goal between group behavior models and individual models based on efficiency for the agents that would have to move around the barrier to exit with the group. While previous research has shown how these groups should move, motivations for doing so have not been widely explored in models (Silveira, Prestes, and Nedel 2008).

1.3 Limitations and Necessary Improvements of Simulation

Despite future computing developments, there will be things that simulation can never accomplish. It may be possible to add basic behavioral elements, but attempting to simulate free will creates obstacles that cannot be mitigated by creating more rules. The human experience is one of infinite possibilities, making exact models impossible to create (Rucker 2004). Furthermore, there are infinite numbers of low probability events that may be appropriate in models, but it is not likely that accounting for them is a good use of a model builder's time. There are always going to be marginal issues that will have to be included or excluded at the discretion of model builders. The infinite possibilities of any real world scenario in a way free model builders. There is often a discussion of positive versus negative modeling, and criticisms frequently focus on elements that are left out of a model (Aguirre, El-Tawil, et al. 2011, Murphy and Salchow 2007). However, with infinite possibilities, whether a model is positive or negative becomes largely irrelevant. Any positive model leaves out an infinite number of possibilities, and any negative model takes away an infinite number. Since it would take an unending amount of time to produce an infinite system, any model of the real world is by definition a compromise, so we can stop striving for perfect accuracy. There are always going to be things that are left out, so model builders should instead focus on what elements are put into a model and why. This discussion necessitates an interdisciplinary approach to model building, which is a requirement for complicated simulations of human behavior.

1.3.1 Communicating Data and Methods

Unfortunately, most evacuation models are proprietary, meaning that the code used to create model output is not available for peer-review. The issue of access to model data and methods compounds the misunderstandings between programmers and model viewers. Many models are considered black boxes by publishers and users, and the peer-reviewed publishing format is not an ideal medium to discuss large amounts of computer code. This makes it difficult to understand exactly what types of reasoning are used in models, and readers are forced to accept what model creators claim at face value. Many models are enormously complicated, and it is not possible to explain methods or parameters in full detail. Additionally, many model packages are expensive or very difficult to learn, and while an interdisciplinary approach might produce better models it is often impossible for most people to contribute to the actual model creation. It is also often difficult to identify the reasoning for avatar behavior. If there are hundreds of things affecting an avatar's “decision”, it is often not a single parameter, but a basket of parameters that work together to achieve the desired objective. Platform specialization is also a problem. Most models only examine a few factors in detail, and continually adding more factors become prohibitive due to issues of programming, cost, and computational power.

However, there are some oversights in specification of human models that are not dismissible as being too difficult to implement. Humans are inherently social creatures, and yet many modeling standards treat human proxies in a manner similar to inanimate objects (Pimenta et al. 2008, Zheng, Zhong, and Liu 2009). This approach produces results through inaccurate reasoning, and additionally excludes entire fields such as sociology, economics, and psychology that are largely devoted to exploring human behavior. These fields have much to contribute to human simulation, and overlooking their theories almost guarantees models that are not as good as they could be. Incorporating social science in future simulations will allow models to be based on more accurate assumptions. In order to encourage adoption of social behavior models, it is vital that model builders discuss methods, data, and types of analysis. This dissertation, which provides high levels of detail about the input and output data of the model used to analyze the Station, is an example of a more open discussion about model results intended to facilitate easier replication.

1.4 Modeling Scenario: The Station Nightclub

This dissertation is intended to push science forward, but it is inappropriate to attempt to do so without acknowledging the sensitivity of data about deadly hazards. While most of this effort revolves around analyzing numbers and iterative improvements, this model exists to attempt to explain a real-world scenario, complete with deaths, injuries, and lasting emotional pain. During the last year of research for this dissertation, John Barylick, an attorney involved in the settlements after the Station fire released a book titled, *Killer Show*, putting the fire and its aftermath in a human context (Barylick 2012).

In the scenario of The Station, there are 465 avatars at the model's inception, when a fire hazard develops in a building during a concert and avatars have four minutes to escape the building alive. About 400 of the avatars were in the building as members of various social groups, making the social component in this scenario very important. The Station nightclub was hosting a band, and a majority of the participants in the scenario were in the building as patrons. This created an environment with a high incidence of intimate social group members. Social groups in The Station scenario ranged from spouses, dating partners, friends, and coworkers. Due to overcrowding of the building (according to occupancy limits), there was also high human density in the confined spaces of the building.

The hazard in The Station nightclub scenario was a fire started by coordinated pyrotechnics that were set off in close proximity to dry acoustic foam. This created a fast-burning fire with multiple points of ignition. Due to confined space, ample fuel, and multiple starting points, the fire quickly spread to most of the interior wall of the stage area and thick smoke filled the stage room in under a minute, creating a difficult environment for evacuation (Grosshandler et al. 2005, Webber, Wright, and Cook 2001). High-density situations such as the evacuation of The Station result in unique evacuation environments, oftentimes making it impossible for agents to remain on their desired paths (Schadschneider et al. 2008).

After the fire at The Station, Disaster Research Center (DRC) researchers obtained detailed demographic and interview data of the victims and survivors of the fire (Aguirre, Torres, et al. 2011, Torres 2010). These data consisted of individuals' demographic information, their familiarity with The Station, their location at the start of the fire, paths used for exit from the fire, observations about social relationships, and individual outcomes. After analyzing these data, it was clear that social relationships played an important role in this evacuation.

One unique aspect of the scenario of The Station was widespread group member dispersal within the building at the time of the fire ignition. The fire began early in the evening, and many social groups had members in opposite ends of the building because the main bar was located far away from the stage. At the moment of the ignition of the hazard, many of the most intimate social groups had all of their members located near one of the various exits, but far away from other members of their respective groups. While most evacuation models would discount group membership and have patrons quickly leave the closest exit, the data collected by Aguirre et al (2011) showed that this was not what occurred. Instead of exiting immediately, many group members attempted to find other members of their social group.

Previous models of The Station nightclub fire focused on either the physics of the fire itself (Grosshandler et al. 2005), homogenous individual evacuation (Pan 2010), or network models (Spearpoint 2012). These models are well thought out and executed, and each does a good job at its stated purpose. However, these models do not include complex social group dynamics, crowd density issues, leadership dynamics, real-time visualization ability, or detailed analyses of individual behaviors and outcomes.

1.5 Dissertation Goals

This dissertation provides a modeling platform designed from an interdisciplinary perspective to be used by social scientists, economists, and computer scientists working together to improve social science and economic experimentation through simulation. There are five deliverables discussed in the dissertation:

1. Detailed results and analysis of the model of The Station with individual and group social behaviors and crowd density effects;

2. A conceptual map for creation of a generalized platform to use the code developed for the Station scenario for other evacuation scenarios;

3. A discussion of validation methods used and created to verify model accuracy;

4. A discussion of the improved model visualization methods created for the models;

5. A discussion of changes necessary to data sets in order to use and share location and relationship data in a way that preserves subject privacy to current academic standards for future applications and generalized models.

1.5.1 Detailed Results Comparison of Model Generations

The primary objective of the dissertation will be the collection and comparison of three layers of model results from the second-generation model developed with social-science principles in place. The three layers compared will be level 1: overall survival and death, level 2: frequency of exits used and overall death, and level 3:

individual tracking and motivation of each avatar. The levels 2 and 3 make the standards of accuracy suggested in this dissertation possible. Instead of benchmarking to previous models or nearest exit guidelines, the model results are compared to the detailed real world results.

1.5.2 Generalized Platform for Group Behavior and Crowd Models

The dissertation returns an example generalized platform that the disaster research community can use to create models that include social behavior elements and crowd density as primary factors, enabling the testing of social science and economic theories. This generalized platform, developed with the lessons learned from the creation of models of the evacuation of The Station nightclub allows the academic community to validate and improve aspects of the model code. Additionally, this platform includes all levels of data tracking developed in the model, allowing for unprecedented levels of model data analysis for future scenarios.

This model addresses some of the unique challenges related to modeling large and multistory buildings (Fang et al. 2012, Galea, Sharp, and Lawrence 2008, Xu and Song 2009). To accomplish this the platform integrates some software conventions from CAD and GIS programs to improve the number of experienced users (de Silva and Eglese 2000).

1.5.3 Improved Validation Methods

Validation methods developed to determine accuracy of models when measuring avatars as individuals are made available to the community for analysis and criticism. The class of models developed for the dissertation includes novel methods of model data collection, suggested for all future models. New methods of validation are required because traditional methods of simple survival counts do not give theorists the information needed to determine why models generated specific output. There is a significant literature devoted to model validation and simulation accuracy, but avatars are rarely analyzed as unique beings (Balci 1997, Kleijnen 1995, Mihram 1972, Olsson and Regan 2001, Sagun, Bouchlaghem, and Anumba 2011, Sargent 2008, Shih, Lin, and Yang 2000, Youngblood and Pace 1995). Because this class of model tracks individual agent paths, goals, motivations, and decisions, a number of theories are testable with the model. The availability of this information requires new validation methods to determine what constitutes a good fit.

1.5.4 Improved Visualization Methods

As a part of this dissertation, model visualization methods developed to better communicate model functioning are made available to the community for analysis and criticism. These visualization methods allow those without a background in statistics to construct simulations and watch models in a visual manner to test theories and examine why a model is functioning in a specific way. Because users can control all aspects of model playback and information output, it is possible to analyze models without ever leaving the projection screen. Additionally, avatar visualization within the model is a topic of interest in the dissertation, since thick smoke was a factor at The Station.

1.5.5 Improved Data Collection and Analysis Methods for Future Study

This dissertation is possible because of the use of detailed network data, tied to places and times. The model validation process has shown that current standards of data anonymization techniques are lacking if relationship data and demographic data are available. This dissertation discusses in the conclusion and appendices enhanced anonymization methods necessary to use location and relationship data in a way that truly preserves user identity while maintaining the ability to apply my model to a larger set of instances.

1.6 Dissertation Outline

These five goals are discussed in six dissertation chapters covering model development, model optimization, model results, and future work. Chapters 2 discusses the development of the first-generation modeling attempt. Chapter 3 discusses the differences between the first and second model, why these changes are significant, and the ultimate purpose of the second-generation model. This chapter includes changes to data collection from the model and changes in model visualization designed to make the model useable by the general population. Chapter 4 covers the development of the second-generation model, used for the rest of the dissertation. Chapter 5 discusses the results of The Station simulation in detail, including industry-standard output validation methods and methods of my own design built to question the importance of clearance time in conventional models. Chapter 6 discusses limitations in the research, ideas for the future, and guidelines for data collection making it possible to model multiple future events in incredible levels of detail.

These six chapters combine to form an original contribution to modeling research by incorporating social science theories about social interaction, from the unique point of disaster science and management. While there are no claims that this model is accurate for any scenario besides the Station, it is designed to allow any necessary future modifications, and released knowing that improvements will be required over the years. It is my hope that this research contributes to the increased adoption of social behavior modeling in building evacuation research.

Chapter 2

PRELIMINARY MODEL: FIRST-GENERATION SOCIAL BEHAVIOR MODEL DEVELOPMENT

This chapter describes the design and development of the first-generation model of the evacuation of the Station using social and crowd behavior elements, and discusses the preliminary results of the study, which inspired the rest of this dissertation. This chapter includes a discussion of four main modeling tasks: individual behavior development, heterogeneous avatars and environment, group behavior, and the behavior of collectivities. A discussion of the results of this first-generation model concludes this chapter.

2.1 Individual behavior model – Task 1

The modeling objective initially had a few relatively specialized requirements. An accurate floor plan of a building was required, so any platform would have to allow for specialized simulations of physical environments in at least two dimensions. The model required simulation of one or more hazards, meaning that several types of agent classes and interactions were going to be required. Finally, complex levels of interaction would be required among the agents because the model replicated humans. The incident modeled had hundreds of people who were together in a relatively small area and duration for a short time period, a situation that would create a huge number of interaction variables when considering individual, group, and mass behaviors.

There are a number of entry-level packages of agent-based modeling platforms like NetLogo, but they sacrifice depth in favor of accessibility (Wilensky 2011). While NetLogo and similar packages offer a friendly and accessible platform to begin agent-based modeling, there are no tools available to extensively modify their environments and create heterogeneous agent behavior. Since the model in this work would not require constant end-user input, a more complex system would be a better fit.

The next layer of modeling platforms included more complicated systems like SWARM in Objective C and Java, and Repast Simphony in Java (Argonne 2011, Swarm 2011). Repast appeared to be a promising platform, and was the basis for the first prototype. While programs like Repast Simphony create powerful agent-based models, components to simulate and test social science theory are not included in the platform at this time. One promising program, buildingEXODUS, is often used for building evacuation research, but the program is proprietary (Gwynne et al. 2001, Owen, Galea, and Lawrence 1996).

After evaluating commercial options, the DRC chose to use EgressSIM for the first-generation model. This platform was chosen because collaborators Dr. Sherif El-Tawil and Mr. Vladamir Federov previously used it to create a model of a building evacuation that had a 2D environment, effective parallel processing ability, a graphic output, and the ability to create heterogeneous agents, something that would have taken a long time to duplicate correctly.

2.1.1 Model Specification

After testing various modifications to the EgressSIM platform, the DRC created an accurate blueprint of The Station at the time of evacuation using a floor map of the building published by NIST, security footage, and interview accounts (Grosshandler et al. 2005, Aguirre, Torres, et al. 2011). This involved mapping a set of exterior and interior walls, doors, exits, windows, and interior obstructions. This map was then converted into a two dimensional environment in EgressSIM. Then we used participant data and determined the location of all employees and patrons at The Station at the beginning of the incident, to build profiles for them in the program (Torres 2010). EgressSIM currently requires agents emulating people to have line of sight to at least one exit. Luckily, this is similar to code requirements for exit signs, so checkpoints exist in the model at similar locations. With these steps completed, the result was a primitive model of The Station using the EgressSIM platform.

As received, the EgressSIM model had one human agent behavior profile, single choices of sequential exits for each agent, no hazards, and virtually no crowding. The model came with 100 identical human agents in a large building. While the initial model displayed a relatively simple physical environment, the design was both easy to understand and efficient, making it an ideal platform for modification and expansion.

Since the model involved evacuation from hazards, additional environmental classes were added to EgressSIM. Windows, which are a hybrid of walls and exits, allow human agents to move through them under duress, but treat them as barriers in a non-emergency environment. Fire and smoke, which damage human agents, each affect agents uniquely. Fire should damage agents, and this effect should decrease with distance. Smoke should affect an entire area more consistently.

The avatars or agents have to compensate for multiple exit paths, changing physical environments, crowding and other limitations caused by other agents, and their own behavioral changes. The most significant obstacles to exit in the model are other agents, an enormous change from the initial environment. Recreating the events at The Station is much more complicated than an orderly building evacuation. The nightclub was overcrowded, and exits were not intuitive or easy enough to reach (Grosshandler et al. 2005). The fast spread of the hazard created a very dangerous environment that changed very quickly, making primary exit routes impractical or impossible to use in a matter of seconds. Having hundreds of human agents negotiating this environment is a difficult task, but by refining behaviors, more realistic simulation became possible.

To be historically accurate, the model should have a finite run time of less than five minutes where human agents could survive without evacuating, a fundamental change from the initial EgressSIM architecture. Because the environment quickly became too hostile for survival for all persons, the model has a finite run time of four minutes. The initial EgressSIM model did not stop until all agents had evacuated the environment, but this is not a good indicator of success in environments where life is at risk because of a hazard.

The static environment in EgressSIM was modified to allow for physical environmental changes based on time or other measures of deterioration. For instance, interview analysis showed that various exits became unavailable during the evacuation, and the model replicates these environmental changes (Aguirre, Torres, et al. 2011). These changes force agents to reevaluate their exit paths, even if the exits that became impassable were not their current targets. Adding more complexity, the ultimate focus of the model is group behavior, which adds additional layers of agent interaction and complexity. Researchers previously identified groups and their degrees of intimacy, and the model uses formulas that determine how likely and for how long members of these groups are to evacuate together or search for each other. Similarly, there are conditions in the program where group members may revert to individual behavior, creating a very complex dynamic for other group members.

2.1.2 Model Creation

Model development began with the creation and modification of a number of EgressSIM items.

Walls: A wall is an impenetrable barrier. An avatar cannot see beyond a wall, and cannot move through it or occupy the same space that is occupied by a wall (for instance, a torso or shoulder cannot be in the same space at any time). If this behavior is not contradicting a goal, an agent will move along a wall as opposed to another path. A wall is a member of the Wall superclass.

Windows: A window is a barrier that is penetrable if necessary. In our model windows do not open, they have to be broken. Avatars see windows as barriers much like walls, except in situations where they are gravely threatened, at which point they will see them as exit or destination options. A window is a member of the Exit superclass.

Checkpoints: A checkpoint is a sign visible to agents, such as the “exit” signs that are required to be visible from most locations in commercial buildings. An agent can see it, and can choose to move towards it as a goal or not. However, any avatar that reaches a checkpoint will then adopt another goal. A checkpoint is a member of the Exit superclass.

Doors: A door is similar to a checkpoint, but it represents a physical item. A door has protrusions on each side, giving similar viewing dimensions to real doors. The protrusions also make sure that walls or other obstacles do not occupy the same space. A door is a member of the Exit superclass.

Destinations: A destination is a barrier that when crossed, removes an avatar from active status in a simulation. These are proxies of exit doors, and assume that any avatar that crosses them is now safe from hazards in the model. A destination is a member of the Exit superclass.

Fire: Fire is an element that damages an avatar based on the distance between the fire and the avatar. Avatars are damaged by fire at a rate that depends on the square of distance (for example, being one foot from a fire is four times as dangerous and being two feet from it). Fire is enormously damaging at close distances, but this threat quickly diminishes with distance. As long as the behavior does not conflict with a higher goal, an avatar will move away from fire. Like a wall, a fire is visually and physically impenetrable in our model. A fire is a member of the Fire superclass.

Smoke: Smoke is an element that damages an avatar largely independent of avatar location. This trait may change in the future for modeling very large or multistory buildings, but in a smaller integrated environment, it is adequate. Because damage is indiscriminate, an avatar does not alter any exit goals because of smoke. Smoke is a member of the Smoke superclass.

2.2 Creating a Heterogeneous Model – Task 2

Models that assume that all participants are homogenous are essentially flow models, measuring human evacuations in a similar way to which sand flows through an hourglass without including motivations (Bonabeau 2002). In task 2 of the model, the focus is the heterogeneous traits of individual avatars. Instead of relying only on fixed parameters such as the closest exit or individual programmed parameters, the model accounts for individual demographics, preferences, and environmental changes caused by other avatars. The model is dynamic and the environment changes based on the position, action, and intention of each avatar.

The individual layer begins like other simulations, with basic physical traits, such as physical size, forward and lateral speeds, and energy values. From there, the model adds heterogeneous traits of age, gender, familiarity with The Station, and willingness to crowd or overtake other agents. These factors create a population of unique avatars.

Specifically missing from the model are data that would bias the model to particular results, such as the ability to survive.

This model only uses avatar input data known before the hazard took place. This prevents the creation of overly deterministic models by using data from the actual results (such as survival percentage or exits used) to influence the model during a simulation. Including these values would create a model where the results are heavily biased. Data from actual results are not part of avatar profiles. The only use of these data are to compare the model results to the historical event.

Based on demographic traits that were identified by the DRC, different profiles for human agents based on age, gender, location, familiarity with The Station, and group traits such as size of groups and group type were developed (Aguirre, Torres, et al. 2011). The simplest elements are those that are external to the model. Individual base behavior calculations come from familiarity with the Station, age, and gender. Familiarity with the Station would imply that avatars would know the layout of the building, including obstacles, dead-ends, and exits. Age and gender affect many behavior variables in the model.

The next type of variable can change based on the environment affecting each agent. Forward speed is a current value and has a maximum value for each avatar. These are the values in meters per second of how quickly an avatar can move from one location to another in the environment; a maximum running speed in an uncertain indoor environment. Lateral speed is the current and maximum speed that an avatar can turn.

This is the value in degrees per second of how quickly an avatar can change direction in an uncertain indoor environment.

"Following distance" measures each avatar's willingness to enter the personal space of other avatars in the environment. This is the minimum amount of space that an avatar will voluntarily leave between themselves and the avatar in front of them (if applicable). Unless density forces avatars closer than their preferences, avatars will not crowd each other beyond their individual preferred following distances.

Avatars that are unfamiliar with the building are more likely to follow interior walls to an exit if an exit path is not abundantly clear to them. If an avatar is unsure of where to go, they may walk until they discover a barrier and then follow it until they come to a point where they have to make a choice about which path to take. Avatars that are more familiar with the environment may be more willing to take paths that are more direct to goals. If an avatar attempts to get to a goal that is not in line of sight, it would be much more willing to take the shortest direct path.

The willingness for avatars to modify their choices is very important to the model. If changing circumstances do not make avatars reconsider their options, then there would not be variation in the model other than randomness. However, in our model, location, time, hazards, other avatars, groups, and the density of the gathering impact avatar choices, and as such, avatars have to be willing to change their goals. In this model, all of these parameters are changing for each avatar over time, and their base value influences their willingness to change.

There is a base value of each avatar's willingness to overtake another avatar. An avatar with a high willingness to overtake would have no problem diverting their path to go around something in their way. An avatar with a low willingness to overtake would be more willing to wait behind others in the way of their goal instead of choosing another path. Again, other factors in the model may override this base value.

2.2.1 Hazard Components

The model of The Station centers on a hazard, a fire that occurred in a confined space necessitating an evacuation. The first consideration was whether to create a hard-coded hazard profile, or run a simulation of a fire during each simulation. Because the NIST technical reports on this fire and how it spread in the building were available, a hard-coded fire model was selected (Grosshandler et al. 2005). However, the framework exists in a way that allows for changes of that component to the simulation for later experiments with different fires or different environments.

Originally, the hazard had two components: fire and smoke. Fire affected individuals at an inverse square of distance, and smoke affected everyone in the environment identically. This was an oversimplification based on a review of fire science models (Cheng and Hadjisophocleous 2011, Miltiadis 2010). Specifically, a detailed report about autopsy and survival data from a medical examiner at The Station implied that the existence or lack of a number of gasses in the environment caused death (Gill et al. 2010).

The hazard portion of the model focuses on the following factors: fire, decreased visibility due to smoke, temperature, oxygen percentage, and carbon monoxide percentage for each avatar. Again, because of the layout of The Station and the propagation of the hazard are known, many of these factors are hard-coded to remain the same in simulation runs, since at this time the model is not concerned with the hazard itself but with the human behavior in reaction to it. Every hazard factor with the exception of visibility from smoke has a point reduction value in each location and time (for example, -1/second). Visibility impacts avatars by eliminating some movement choices and communication types (Isobe, Helbing, and Nagatani 2004). Each avatar has a resistance value that these points are subtracted from, calculated from a combination of demographic factors. Based on the data gathered about The Station, and the reports from the incident, we know generally where participants were at the start of the fire, how they exited or perished, and what their main cause of injury or death would have been based on their probable locations (Aguirre, Torres, et al. 2011). The desire is

to produce results that are accurate at any point in the simulation, not just at the end (measuring the correct number exited or deceased).

2.2.2 Single Agent Behaviors

Next, there are several behaviors of a single agent in the environment. In the environment, a single avatar that is alone can use the following rules:

Behavior Type: The behavior type is a set of attributes that an avatar possesses individually.

Energy Level: The energy level is the initial allocation of “health” for an avatar. Fire and smoke damage deplete the energy level.

Initial Location: The initial location is where the agent is located on the (X,Y) plane when the model begins a simulation run, corresponding with the beginning of the fire. This location must be visible to at least one member of the Exit superclass, and must not overlap with a member of the Wall or Fire superclasses.

Lateral Movement: Lateral movement is the avatar's ability to turn, dictated by Behavior Type and barriers. The axis for lateral movement is the center of the torso, which does not change location.

Forward Movement: Forward movement is the avatar's ability to perform the equivalent of a step, dictated by Behavior Type and barriers. Forward movement changes the location of the avatar's center.

Target Point: The target point is the location that an avatar attempts to move toward to accomplish a goal.

2.2.3 Moving in the Environment

A single avatar initializes in an environment and evaluates available potential Target Points (which all possess a potential value, for instance a destination has a higher potential than a checkpoint). The Target Point potentials are descending, the lower a value is, the better. When a point is selected based on availability and distance (with the lowest potential value), the avatar will attempt to move toward it, using lateral movement and forward movement, avoiding fire, trending towards walls if applicable, and reevaluating their movement at every step.

Goals may change before an agent reaches a Target Point, but they mandatorily change afterwards. If an avatar reaches a Target Point and the avatar is still active in the simulation, they will reevaluate and find a new Target Point with a lower potential value than the point reached previously. If one is not available, the agent will return to the previous point and reevaluate. These calculations occur 20 times a second.

2.2.4 Adding Avatars

Adding even a second avatar creates a great deal of complexity, and these issues become a greater concern as density of avatars increases. Based on each avatar's behavior model, there are traits that account for comfortable space between avatars, willingness to bypass other avatars, willingness to change goals, speed, and willingness to bump into other avatars. Many of these traits are only active on occasions where avatars are crowded in with others.

A unique aspect of our model is the ability for individual behavior models to change based on environment. As hazards become more pronounced, or current parameters are not working, avatars should be willing to change their behavior to include traits that would not exist in a normal situation.

Crowding and density are important to the simulation, and it is important that the behavior profiles are accurate at every level. Goal and movement calculations become much more complicated when many of the barriers are fluid (other agents). Through no fault of the individual avatar, it is possible to select goals that are nearly impossible to obtain because of the goals of other agents. If an avatar decides that due to their profile, environment, and other agents, they want to move in an opposite direction of the collectivity, they very likely affect not only their own goal, but the goals of several other agents.

It is difficult to predict the behavior of a group of individuals, so goals are much more likely to change upon recalculation when exit opportunities are blocked or might be blocked by other agents.

2.3 Incorporating Group Behavior – Task 3

The next layer of the model is the group or social layer and resulting social relationships. There are multiple examples of evacuation models incorporating basic social behavior, but their focus is not on relationships (Gwynne, Galea, and Lawrence 2006, Pan et al. 2006). Studies that account for different social relationships in evacuations do not extend to model development (Johnson, Feinberg, and Johnston 1994). This study accounts for multiple different types of social relationships. Social links between the avatars were determined for the Station model. The resulting dataset contained links between avatars, and the type of linking relationship. There are five types of relationships: co-workers, friends, dating partners, spouses/family members, and multiple relationships (avatars that have more than one relationship type with other group members). The different types of relationships have unique features in the model. For example, co-workers are much more likely to revert to individual behavior than spouses or family members.

2.3.1 Leadership

Following previous ABM works, the model puts group members into two main categories, leaders and followers (Pelechano and Badler 2006). Group leaders are at a given moment in charge of their groups, meaning the leader has determined which exit path the group will take. All other members of the group are followers, and either heed the instructions of the leader or remove themselves from the social group.

The group member most likely to be a group leader is in charge of the group unless they exit the building, die, split from the group, are overtaken by another group member, or the group dissolves. If the group does not dissolve but the leader becomes inactive, the other group members compare leadership values and the avatar with the highest remaining value will take over leadership of the group.

In the current model, the group selects a leader, and the leader exhibits their individual behavior. The followers in the group then follow the leader if in line of sight, or the leader's target point, if other group members are in line of sight. If a follower loses line of sight of all group members, they will revert to individual behavior after a stated amount of time. If a leader chooses to revert to individual behavior, or leaves the group through external factors (death, reaching a destination) then a new leader will be selected by the group and assume responsibility for selecting target points. If all members of a group but one are inactive in the group, the remaining member will automatically return to individual behavior.

The group leader is chosen by an equation that gives preference to “middle” age (in the Station environment this age is in the high 30s), male gender, familiarity with the scenario, smaller distance from other group members, smaller distance from exit points, and lack of injury. These demographic characteristics were selected based on the specifics of the Station scenario and will almost certainly change when programming for other environments. Leadership is a dynamic factor based in part on the changes brought about by other avatars not in the group and the developing fire and smoke hazards. Based on several changing events, the group leader may also change. However once a leader is chosen there is a premium added to the leader's leadership value. This premium means that a situation will need to change significantly in order for a group leader to change. Without a leader premium, group leadership could potentially change every timestep, resulting in groups that did not move.

Some of these variables are static and do not change during a simulation. Age, gender, and familiarity with the building are all static values that do not change based on what is occurring in the simulation. In the model the male gender, older age, and more environmental familiarity all increase the probability of an avatar being a group leader. Male and older avatars have a greater probability of becoming leaders because of social science and psychology research showing that males and older group members often take leadership roles in emergencies (Kent and Moss 1994, Tueretgen, Unsal, and Erdem 2008). Increased familiarity with the Station building raises the probability of leadership because group members familiar with an environment will be more comfortable in an emergency situation and will have a greater understanding of evacuation options (Mawson 2005).

Once a leader is chosen and the leader premium is added, it is unlikely that leadership will change unless a leader leaves the environment, is unable to move freely, becomes inactive, or the situation in the environment changes drastically. For instance, if an exit that the leader was leading the group to were to become inaccessible, it would be more likely that group leadership would change, since the leader would now likely be the furthest from another exit. The leader premium is important because without it, the leader values might change too frequently. A leader premium prevents leadership change while an exit plan is in effect while still allowing for leadership changes that should occur due to the changing environmental circumstances.

2.3.2 Group Behavior Specification

Evaluating group behavior of agents creates an additional level of complexity. After adding social behaviors, the issues and interactions raised with multiple agents become even more complicated. There is an entire additional set of modeling complexities that are raised when attempting to coordinate groups of avatars in a physical environment (Aguirre, Torres, and Gill 2009). Rules of timing, space, and movement are layered upon communication protocols and translators required for group networking.

Group size and group intimacy both contribute to the importance of groups to each avatar.

Currently, the model uses two group behavior layers, essentially outlined as a follower layer and a leader layer. The eventual goal is to create groups with possibilities for decentralized leadership formats (such as group decisions and negotiations), but for now, there is a dynamic leadership component.

Of the 465 avatars the DRC modeled for the Station scenario, only 48 were alone. The other 417 avatars (90% of the total) are members of social groups. These groups break down as follows: There are 68 groups of two (136 avatars), 25 groups of three (75 avatars), 19 groups of 4 (76 avatars), 9 groups of 5 (45 avatars), and 85 avatars in groups of six or more. 44 avatars are defined as "alone" (no group association), 54 avatars defined as co-workers, 193 as friends, 57 as dating partners, 40 as spouses or family, and 77 as "multiple" (for instance friends and partners).

The first group layer consists of social structure. There are different parameters for intimate groups and social groups. If an avatar is a member of a group, these traits will override or complement their individual behavior models. Instead, they share goals and movement with their group, or at least attempt to do so. However, if it is impossible to follow the group goals, or there is an imminent danger detected, an avatar in a group maintains the ability to revert to individual behavior.

The social structure layer highlights the social relationships between avatars and resulting behavior changes compared to individual motives. In the model, there are five different types of group relationships; co-workers, casual friends, dating partners, spouses and family members, and those with multiple relationship types.

The co-worker group type is the loosest group association in the model. Co-worker relationships are often well-defined but also very casual (DeVito 2003). In the model, the definition of a co-worker is someone that an avatar knows, but would not choose to associate with outside of work. Because of this definition, most of the co-worker relationships in the model are employees of the building. Because of the formal and forced aspects of co-worker relationships, the co-worker bonds in the model are not strong. This means that agents with co-worker bonds will be less likely to avoid individually optimal exits to instead search for co-workers, and are less likely to follow other co-workers on paths that do not appear to be individually optimal. The loose co-worker bonds are likely to disintegrate early on in the presence of a hazard.

Casual friends are another loose group association. In the model, casual friendship or friendship is defined as a relationship that is maintained in free time or because of a common activity (DeVito 2003). Many of the avatars in The Station scenario meet this definition, because there were several groups of casual friends attending the nightclub. Casual friends are more likely to search for each other or stay together than co-workers, although they will not abandon clearly optimal exit choices to do so. When the environment begins to deteriorate, these group members are likely to revert to individual behavior.

Dating partners are unmarried couples with a more cohesive group association. Dating partners are for the most part in intimate groups of two, which makes it much more obvious if an avatar's partner is missing (DeVito 2003). Dating partners are more likely than other casual forms of relationships to search for their counterpart and to follow a group leader even if the behavior is not individually optimal. Even if dating partners are far away from each other and in clear view of individually optimal exits, they are likely to search for each other despite devotion to self. The environmental situation has to deteriorate significantly in order for dating partners to abandon their group, but for the most part dating partners will revert to individual behavior if they are significantly injured.

Spouses and family members are the strongest group bonds in the model (DeVito 2003). Spouses and family members are almost certain to take time to look for their group members, and are very reluctant to abandon them in all but the most hazardous situations. Spouses and family members make up very cohesive groups where leadership values are strong, groups are likely to converge and then stay together, and group members are unlikely to leave the group unless not doing so will be fatal. Members of these groups are extremely unlikely to evacuate a hazard situation without their other group members.

Avatars with multiple group-level relationships are likely to adopt the strongest relationship of another avatar in the group. If a group member is a co-worker and a friend to other group members, the group member will act as a friend. If the group member is a friend and a family member to other group members, the avatar will act as a family member. This assumes that if the group is together, the member creating each avatar's strongest bond will be a part of the group.

The interactions between avatars and their group members, other agents, and their environment create a complex and fluid simulation environment, where there are tens of thousands of calculations per second. When decentralized groups are used, these interactions will grow, and avoiding or managing conflicting commands will become more of an issue.

2.4 Incorporating Crowd and Reactive Behavior – Task 4

The final layer of the model is the density layer or supra force. The supra force layer depends on the environment and density of other avatars in a specific setting. In very dense situations, avatars have reduced or eliminated ability to move freely; avatars may find themselves surrounded and unable to move in a direction different from those surrounding them. It is important to note that supra force behavior is different from panic behavior, since group incidence of panic behavior was not apparent in the evacuation of the Station (for a panic example see Helbing, Farkas, and Vicsek 2000, non-panic example see Gill et al. 2010). Supra force is an interaction between a collectivity of people and an environment. It is an ecological pattern that is in constant change as the collectivity moves over time. It is also an example of emergent collective behavior that is observed even if unplanned in the results of the simulation model. Supra force is modeled more similarly to crowds watching out for each other even if they cannot move where they choose (Drury, Cocking, and Reicher 2009). If the conditions of supra force are present for an avatar, individual and social behavior elements will be overridden, and the avatar will move in the direction of the movement of the avatars surrounding it and its group. In effect, the supra force condition temporarily supercedes the avatar's desired goals.

2.4.1 Supra Force

The Station scenario is a situation with a high human density in a closed building environment. The scenario involved a building that was significantly over municipal occupancy limits at the time of the hazard (Grosshandler et al. 2005). The Station only had four exit doors, and there were 465 people in the building at the time of the hazard, resulting in a high number of evacuees for each available exit. This scenario results in significant crowding at the paths to the exit doors.

The model proposes that in situations with high human density combined with otherwise inflexible environments, avatars surrounded by multiple other avatars or immovable obstacles in all directions will no longer possess free will to move, and may actually be moved or crushed based on the force of the moving collectivity of the avatars when a precipitating event occurs. Supra force is a condition that can occur during situations with a combination of high density, diverse population movement (often contraflows), and unique environments. It is important to mention that density alone does not create supra force. Building upon previous efforts this model regards collective crowds as discrete, emerging features of emergency evacuation situations resulting from high densities of avatars in a conducive environment (Low 2000, Silveira, Prestes, and Nedel 2008). In the model, this phenomenon is called supra force to emphasize the emergent features of these collective patterns. Avatars that are not able to move in any direction, even if they push avatars next to them experience supra-crowd force, and their behavior model reverts to a passive state. If avatars surrounding the agents subject to supra-crowd force move in a consistent direction, the surrounded avatars may not be able to move in any other direction (Bohannon 2005). The surrounded avatars will then wait until they are able to move to resume their goal. In the group setting, if a group leader is a victim of supra force, a new leader is selected.

In the model, there are specific requirements for experiencing supra-crowd force. The first requirement is that an avatar is surrounded on all sides by an immovable environmental barrier or other avatars. The second is that two or more avatars are present in all directions and surround any avatar not limited by an environmental barrier. The assumption is that an avatar only blocked by one layer of other avatars will be able to push them out of their current position. The third requirement is that surrounding conditions should last for five or more timesteps (or a quarter of a second). This prevents avatars from surrendering their will based on fleeting circumstances.

Once an avatar experiences supra force, they remain in the behavior mode until a second after the situation changes (the lagged time accounts for a reorientation). After the situation that caused the supra force recedes, the agent reverts to their previous behavior mode, although they are almost certainly less likely to transition back into a group leadership role. Avatars may experience supra force repeatedly, as often as the required conditions exist. Avatars that avoid situations with high human density are very unlikely to encounter supra force.

2.5 First-generation Model Results

Preliminary model results are very promising. Aguirre et al. 2011 discusses some of these findings, but an expansion of these results is warranted here. These four layers work together to create a dynamic model where the environment, individual traits, group traits, and density all affect evacuation choices. Put together, these different considerations result in a model that is much more lifelike than a simple closest-exit algorithm. In this model, avatars consider their group members, various hazards, population densities, and improvised exits when choosing the optimal path of exit. Optimal choices for each agent change constantly based on the specific conditions in the model, allowing avatars to make more informed choices based off of new information.

This multiple-layer model is unique because of the simultaneous addition of heterogeneous individual traits, group-level traits, and supra-crowd force. With these conditions included, the model produces results that are significantly different from versions that only include individual behavior. These traits are important for model builders to consider and include in future models, because human beings do not operate in social vacuums. Model builders should embrace social science research about social behavior and create models that use avatars with traits resembling the studied populations.

More detailed results from the first-generation model below are available from Aguirre et al 2011a on page 426.

Table 2.1: First-Generation Model Results

Exit	First Generation Model	Actual
Bar Exit	113	78
Kitchen Exit	7	17
Main Exit	126	128
Stage Exit	24	24
Bar Window	56	71
Greenhouse Window	2	34
Deceased	137	100
Total	466	452

These first-generation results suggest that adding group and crowd behavior traits may create more accurate and realistic evacuation models, at least for the Station scenario, and generally support the idea that social science concepts should be included in evacuation models. In addition to showing the potential of this class of models, these results inspired me to look deeper into the reasoning used for these models, and encouraged me to build a second-generation social and crowd behavior model.

Chapter 3

LESSONS LEARNED FROM THE FIRST-GENERATION MODEL: HOW TO MAKE A GENERALIZED SECOND-GENERATION PLATFORM

While detailed results analysis of the Station scenario are discussed in chapter 5, it makes sense now to discuss some of the coding and theoretical problems discovered in the first-generation platform and addressed in the second-generation model. The largest difference between the two generations of the models is that the second-generation platform is that it can be used to simulate any floor plan and is designed to be used by researchers that are not as comfortable with computer code. The similarities and differences of these models are important to discuss in detail because they highlight exactly how the models work and why features were added or removed for the final versions of the program.

3.1 Building Environment

In both versions of the model, an identical environment was used to map the Station. Blueprints, obtained from a National Fire Protection Association (NFPA) study, were imported into the model environments (Duval 2006). Measurements in the model are in meters, and are identical or as close as possible to the actual building. Because the models render a floor as a two-dimensional environment, there is no raised platform by the dance floor. The raised platform is rendered as a part of the dance floor, and the partial walls that were generally considered not passable are rendered as walls. In the future, a three dimensional rendering of the building environment might offer more accurate results, but that is outside of the scope of these first social behavior modeling efforts.

As outlined in the NFPA study, the Station is a one-story building with a number of different rooms and six potential exit categories: main exit, bar exit, kitchen exit, stage exit, bar windows, and greenhouse windows. These six exit categories are all rendered in the model environment, and their sizes closely match the available information about the Station. There are 8 possible outcomes for all agents: remaining active but in the building (almost impossible in the environment), inactive in the building, and out of the building via one of the six exits (Torres 2010).

The environment is an important part of the development of the model. In the Station, there were a number of specific environmental issues that made the occasion lethal. These included exits without signage, exits blocked by staff, egress paths that were too narrow for multiple people to navigate, dry acoustic foam, a lack of sprinklers, and loss of visibility. These attributes, combined with the high crowd density and the distribution of the crowd created a situation that very quickly became difficult to navigate. These attributes, along with the specific makeup of the crowd created a lethal situation in an environment where traditional evacuation modeling would not predict it.

As outlined in previous DRC research, there are six main ecologies in the Station (Aguirre, Torres, et al. 2011). These main ecologies are bathroom area, kitchen area, bar area, stage area, sunroom area, and main entrance. The distribution of the population is an important trait in the model, and it was important to confirm that all avatars were correctly located in their ecologies and actual location at the start of the fire. The population distribution was important because of the group relationships. Since many avatars attempted to locate their social groups instead of simply taking the closest of fastest exit, correct positioning is necessary in the model. In order to correctly position each avatar, interview and news media data was used to position avatars, and for the avatars that could not correctly be positioned based on prior research, approximations were made based off reports from other survivors or group members (about three dozen profiles).

3.2 Avatar Traits

Avatars in the model have a number of physical attributes based on the characteristics of victims. These attributes include age, gender, type of social group, membership in social groups, previous visits to the Station, and location at the start of the fire. Importantly, no data that could only have been known at the end of the fire (such as exit used, survival, intended exit, and group member well-being) is included in the model. This data is specifically left out of the model because it would bias the results. Because the intention of this dissertation is to show that social science theory-based models can contribute to more accurate modeling, it is vitally important to bias the results in any way.

The static avatar attributes described above are unchanging in the model. At the start of the simulation, groups remain the same, avatar features remain the same, and importantly, the models are not learning models, meaning that avatars do not get better at navigating the environment over different model runs. Because the intention of these models is to validate the premise that social science theory can contribute to more accurate modeling, a learning model is not the correct tool for the task. In the future, learning models may help create better social science theory tools, but as a first effort, it is important to create models where decision-making ability does not change between runs.

One of the most important avatar variables is forward speed. Forward speed determines how quickly avatars can get out of the building. Generally, the forward walking speed of adults (the population in the Station) is between 0.75 and 1.5 meters per second (Al-Obaidi et al. 2003, Fruin 1987, Oberg, Karsznia, and Oberg 1993). Obviously, running speed is significantly higher than walking speed, but inside of buildings where it is not possible to move very far in a straight line, walking speed is a better approximation of maximum possible speed (Predtechenskii and Milinskii 1978, Riley, Della Croce, and Kerrigan 2001). In the model, there are values for maximum speed under duress (higher than normal walking speed) and normal walking speed. Normal walking speed varies with age and gender (Oberg, Karsznia, and Oberg 1993).

Another important speed value is lateral speed, or turning speed. Lateral speed is a value, measured in degrees per second, that dictates how quickly avatars can change direction. Each avatar has an orientation, or a direction their body is facing at any given point. Lateral changes measure the difference in degrees between two orientations. Orientation and forward speed together create a movement profile for an avatar. While it is possible to easily move 360 degrees in under a second, it is not easy to do so under control, or while moving forward (Riley, Della Croce, and Kerrigan 2001). As a result, maximum later speed values are reduced if an avatar is in forward motion. Forward speed and lateral speed are identical in the two versions of the models. Because these are static maximum values and are generally well-established in the literature, there is no reason to change them further.

In addition to speed values, each avatar has values for how willing they are to overtake other avatars, and how willing they are to crowd other avatars. Unlike the speed values, these values can change significantly based on environment. Avatars that are hurt are more willing to overtake or crowd others, while avatars in environments with low levels of stress default to more relaxed values. In the United States, where the Station incident took place, people generally prefer large amounts of personal space, and very high density situations are avoided unless people are under duress (Beaulieu 2004, Sommer 1959) . When adapting these models for other cultures, it will be important to make changes for different cultural understandings of density.

The two versions of the models have slightly different versions of behavior values for avatars. In the first version of the model, there were 11 behavior values from passive to aggressive, which influenced the previously discussed values for avatars. Age, gender, and other attributes contribute to the behavior model in the first version of the model. In the second version of the model, the reduced number of behavior models are modified based on the condition of each avatar. As avatars become more desperate to leave a hazardous environment, their behavior models adjust to more aggressive settings.

3.3 Group Dynamics

Group dynamics are the most important aspect of these social models, because these dimensions are what make the models significantly different from current evacuation models. Most current models use shortest-distance or shortest-time engines to determine which exits to attempt to use. The models in this dissertation take social relationships into account, which frequently cause avatars to make choices that would be impossible in shortest-time or shortest-distance calculations.

In the first version of the model, group type and willingness to leave the social group was a static calculation assigned to each avatar. In the second version of the model, group type is still a static value, but the willingness to leave a group is calculated based off the specific situation of each avatar. This is a major difference between the two versions of the model, and it makes avatars under great amounts of stress more willing to leave their groups.

3.3.1 Leadership

In both versions of the model, group leadership is an important attribute. In both models, there is one leader at a time per group, and the rest of the group members are followers, meaning they are influenced by the goal of the group leader if they are a part of the group.

In the first version of the model, leaders were hard-coded based on study of the survivor interview data. While this was an intensive process, it does not scale to other situations besides the Station.

For the second version of the model, it was important to create a dynamic leadership equation, where leadership is an attribute defined by conditions in the model. Despite a large body of literature on leadership, there is a surprising lack of literature or study about the quantitative qualities that define leaders. The leadership equation developed for the second version of this model is an admittedly imperfect first attempt to quantify leadership selection in an evacuation environment.

The first leadership equation has a combination of static and dynamic elements, including prior visits to the station (familiarity), gender, age, energy, distance to an exit, distance to other group members, and a past leadership premium. Familiarity, gender, and age are static values that do not change over the model run. Energy, or the lack of injury in the model, is a positive variable. A group member with a high energy value is more likely to be selected a group leader. A closer distance to an exit or closer distance to other group members also makes an avatar more likely to have a high leadership score. The past leadership premium is designed to make leadership relatively stable. This value adds to the leadership score or all avatars that were leaders in the last turn, requiring other avatars to become significantly more desirable leaders before leadership changes. Without this value, leadership might change twenty times per second (the number of turns in the model) making group membership extremely inefficient and not lifelike.

3.4 Supra Force

Supra force, or the inability to move freely based on specific high-density situations with contraflows and environmental limitations, is another attribute that makes these social science theory based models significantly different from other evacuation models available today. In very high-density situations, it often becomes impossible to move, and because of this emergence, individual and group level motivations may be superseded by this force.

In the first version of the model, supra force calculations were based on an avatar being surrounded by an environmental barrier or at least one avatar on all sides. This presented an environment where a surrounded avatar would not be able to move without moving another avatar.

In the second version of the model, where avatars can additively push each other, supra force conditions occur when an avatar is surrounded on all sides by an environmental barrier or two or more avatars. This distinction was made because if there is only one other individual between an avatar and a goal they could be pushed, but two or more avatars blocking the path makes it difficult for a lone avatar to push others out of the way. This change creates a more realistic scenario for supra force because it eliminates the border cases in the first version of the model were avatars are identified as surrounded even though they are able to push away one layer of avatars rather easily.

When an avatar is subjected to supra force in the model, their individual or group goals are temporarily abandoned and they move in the direction of the surrounding gathering if it is moving in one direction, or they do not move at all if the gathering is moving in opposing directions. In both conditions, the avatar is not moving by following their own goals or the goals of their group.

In the future, supra force is likely to be another important area in model development for high-density situations. In high-density situations, individual desires are often overruled by the reality that it is not possible to move in one or more desired directions. It is important to try and quantify this condition. The supra force conditions are easy to calculate and not computationally intensive, but they play an important part in creating the accurate results desired from the social science theory-based models.

3.5 Model Output

While the original intention of the first-generation model was only to test if social-science theory based models could return accurate results, it quickly became clear that the benchmarks used by most modelers were inadequate to decide if a model was accurate or not. The entire premise of the first-generation was to see if social behavior theory could explain why agents take exits that might not be closest or fastest, and in order to do that, individual agents needed to be tracked as unique individuals.

Most building evacuation models only track how many agents escape a hazard situation and how many survive. As mentioned in the first chapter, this is considered “level 1” model tracking in this dissertation. The only thing that level 1 results can track is overall lethality, and these overall exit numbers offer no help to explain why exits were used, what paths were taken, where there were delays, or any other issues.

The level 2 results (generally the most detailed results that other models use) track specific exit paths, but fall short of actually considering agents as individuals. Level 2 results allow modelers to analyze which exits were used, and then compare number of agents in an area versus the number of agents that used an exit. This is a rudimentary way to calculate how popular exits are, but these results still do not track specific agent paths.

Level 3, or tracking agents as individuals, allowed us to answer the question, “can a social science theory based model offer accurate results?” Level 3 tracks agents as specific individuals, making it possible to compare the actual exit used by an avatar compared to the exit used by the historical person each avatar simulates.

The first version of the model outputs enough data to calculate results for levels 1-3. These three levels of data are already detailed enough that was not possible to compare all three levels to other models. Comparing the results to the actual situation, the numbers for all three levels were promising. However, while analyzing the data shown in the results section, it became clear that in order to truly address some of the hypotheses in this dissertation, even more data were required.

In an attempt to address all of these hypotheses, the output files were significantly modified for the second version of the model, and the new platform allows for researchers to analyze their future results using this greater level of data collection. The second-generation model outputs six files for every model run. Essentially, these files output level 1-3 data for the end of the simulation, and a second set of three files that output these three levels of data periodically during each simulation at intervals specified by the model builder.

The first generation model collected overall survival data for Level One, and exit data for levels two and three. The second-generation data sets collect enough static data to confirm that agents are correctly identified, and all of the dynamic data created in the model that is not tracked by measuring the end results. This includes location, goal location, orientation, movement, energy, leadership status, group status, and exit used. This wealth of data makes it possible to analyze where agents intend to move, why they are moving there, where they have been, who they are following, whether they are facing their goals, and other information that makes the detailed analysis in chapter five possible. While this level of data collection goes beyond the original stated goals, it is necessary to address the fundamental research questions of this dissertation. This level of data collection should be considered for all models of building evacuations in crisis contexts.

3.6 Usability

Some of the most important changes in the model relate to usability. The biggest difference between the two model versions is that the second-generation model is built on a general framework, meaning that the engine used to create the results can be modified for different situations. Making it possible to recreate similar models will motivate other researchers to consider social science theory based modeling, and encourage others to become modelers for the first time.

3.6.1 User Friendly Code

The first-generation model was not user friendly, requiring users to know how to code. Thus, the second-generation model includes a number of features designed to make it possible for people that do not know how to code to create and analyze models. While this was not a direct objective of the dissertation, it was important to attempt to do. While the second-generation model still needs some work to be made into an executable file, the model framework at this time can be used with a mouse and keyboard, without users needing to know how to code. While there is still a way to go to make it more user friendly, it is currently one of the more user friendly models available, while still allowing for a high degree of customization.

The second-generation model was designed to look and function like the first-generation model, but created a system where a model could be built even if the model builder does not know how to code. The following subsection provides an example of building a model in graphical user interface (GUI) environment, and describes why this is an important feature.

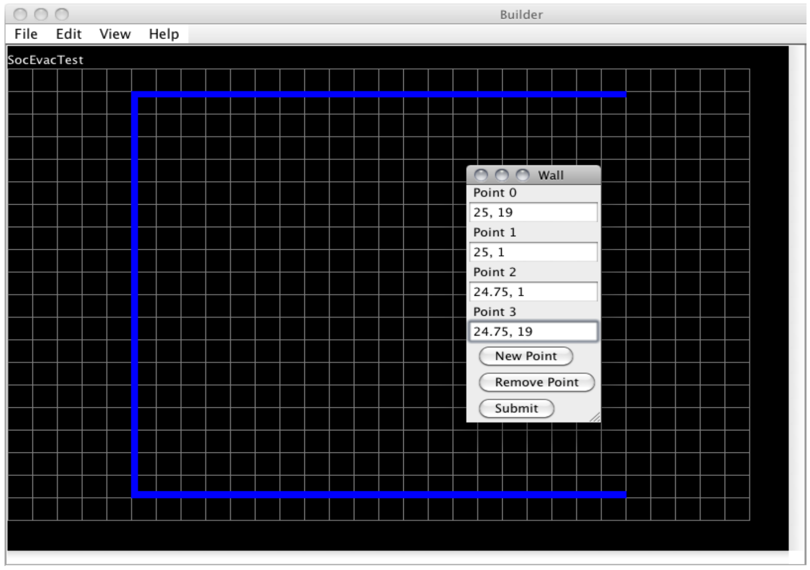

Figure 3.1: GUI Wall Construction

All aspects of the model can be coded, and this is the preferred input method, there is also a GUI available to build model features. Figure 3.1 above shows a wall being constructed using mouse and keyboard input on a GUI.

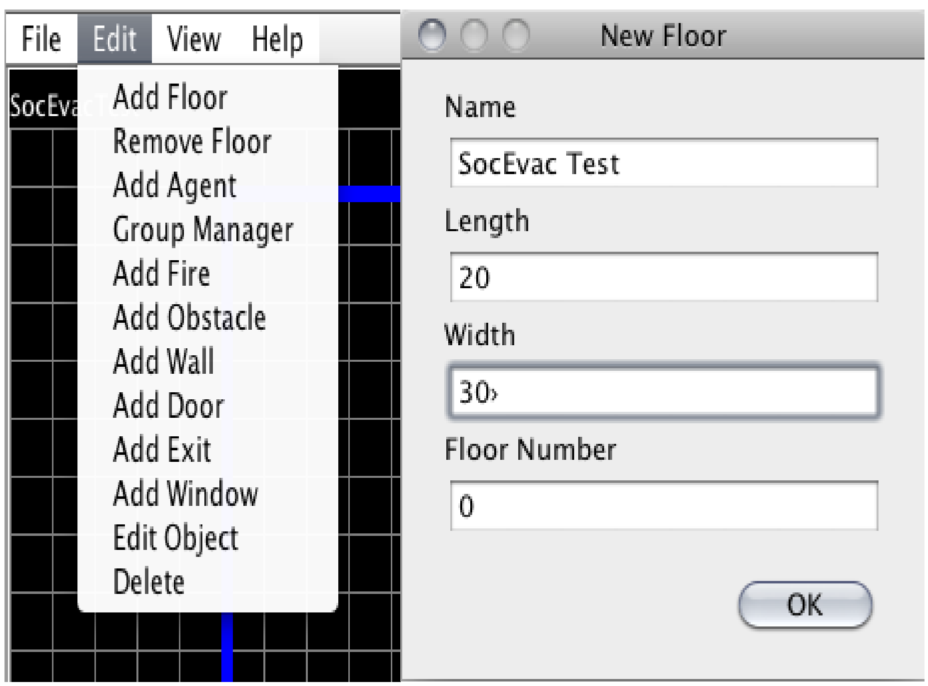

Figure 3.2: Environment Element Selection and Floor Construction

Figure 3.2 shows the process of creating a new building floor in the second-generation model. Using the "Edit" menu, a model builder can select which attribute to construct, which will result in a popup window. The user can then select values with a keyboard, or input certain values by clicking on the floorplan with a mouse. This process allows model users to quickly build or modify models once they are created.

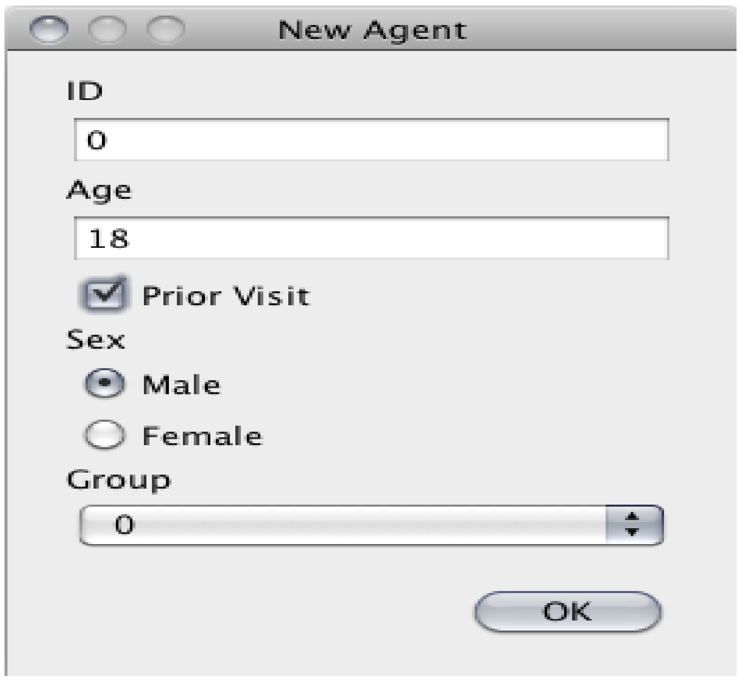

Figure 3.3: Agent Construction

Figure 3.3 above shows the GUI process of creating a new agent in the second-generation model. All of the static values (age, gender, group, prior visit) can be selected at the time of creating the agent, and the user can click the point in the environment where they would like the agent to initialize. This process makes it fast and simple to add agents to the modeling environment.

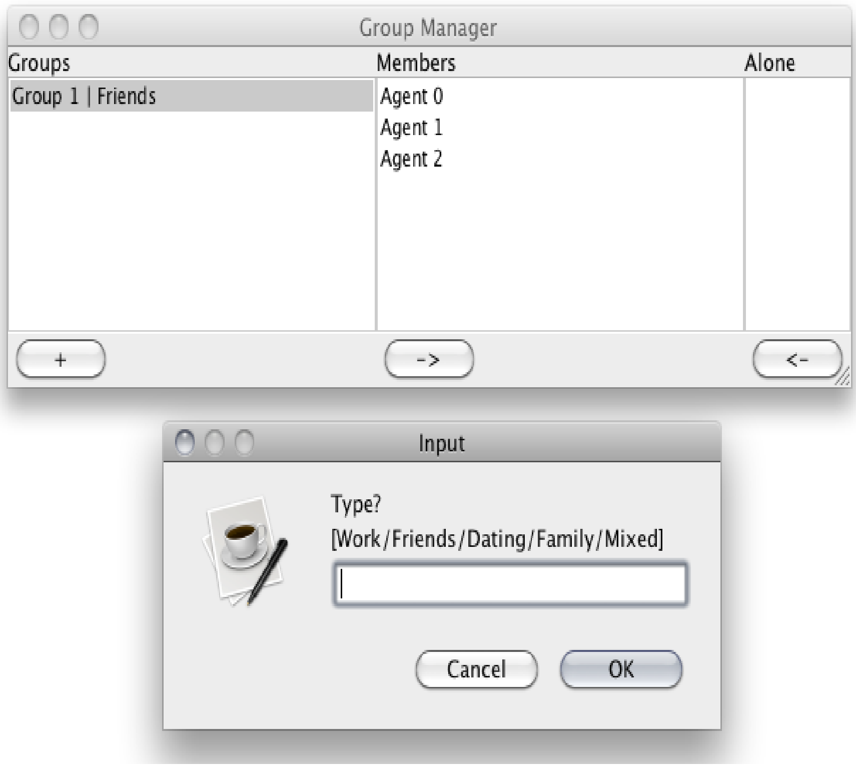

Figure 3.4: Group Builder

Figure 3.4 above displays the GUI windows related to group management, where model users can choose what groups the agents belong to, and how the groups are categorized. These group features are based on the design of the groups in the Station, but it is possible to add other group types in the future. Importantly, the group members can be managed without accessing the profiles of the individual avatars, which would be difficult to do in large simulations.

3.6.2 User Friendly Output

In order to increase user friendliness, the output of the second-generation model was changed to independent named comma separated values (.csv) files. The first generation model output data in text. While this was adequate to analyze the model, the second-generation model format makes it easy for users to look at results, and also makes it easy to run the model many times without intervention to gather output data. The second-generation model makes it possible to run hundreds or thousands of simulations and collect all the data afterwards. Over the course of my dissertation, this became a very important feature making it possible to increase batch size for comparisons.

The six files output in each run of the second-generation model are output as .csv files with a clear name and a timestamp at the beginning of each file name. These files are then saved in the output file of the model directory. The files thankfully have very different sizes, making it possible to sort the files by size, isolate each file type, and then use some simple terminal commands to merge the files into one file. Unfortunately, it is not possible to decrease the file sizes much more than has already occurred, meaning that 1,000 or 10,000 run batches still create files that can be several gigabytes in size, making them difficult to analyze on consumer systems. This is a problem that it is not possible to address without sacrificing data quality. Thankfully, the model framework is small in size, making it possible to run the model on avareage consumer systems in 2012.

	A	B	C	D	E	F	G	H	I	J	K	L
1	Time	AgentID	GroupNumber	Floor	isDead?	EnergyLevel	LocX	LocY	TargetX	TargetY	Following?	Surrounded?
2	0	0	0	Test	FALSE	100	2.75	18	5.25	18	Not implemented	FALSE
3	0	1	0	Test	FALSE	100	5.25	18	6	10.25	Not implemented	FALSE
4	0	2	0	Test	FALSE	100	7.75	18	5.25	18	Not implemented	FALSE
5	0	3	0	Test	FALSE	100	3.25	17	5.25	18	Not implemented	FALSE
6	0	4	0	Test	FALSE	100	4.75	15.75	5.25	18	Not implemented	FALSE
7	0	5	0	Test	FALSE	100	6.75	16	5.25	18	Not implemented	FALSE
8	0	6	0	Test	FALSE	100	8.5	15.25	5.25	18	Not implemented	FALSE

Figure 3.5: Agent Output

Figure 3.5 displays an example of the largest of six files output from every simulation run. This file presents detailed data for all of the individual avatars at different time points specified by the user. The columns visible in this sample output include time, agent, group number, floor (of building), cause of death, energy level, location, location of goal, leadership status, and supra force status. These files are accessible after the simulation has finished running, and the files are directly comparable across simulations, making it possible to compare large batches of files. Real-time visualization of simulation runs is discussed below.

3.6.3 Visualization

There are two ultimate purposes of the current visualization system: to allow users that are not comfortable with statistical output to verify the results of the model, and to allow researchers to confirm in real time that the model is functioning as intended. The selectable visual variables make it possible to confirm that group members are looking for each other or navigating the environment together. The supra force visualization makes it very easy to confirm that avatars that should be completely surrounded are flagged appropriately, and that avatars that can freely move are not incorrectly flagged as experiencing supra force.

These visualization tools are valuable when calibrating the model. A properly implemented visualization system makes it possible to identify problems much more quickly. This is a significant time saving feature that should be incorporated in more models

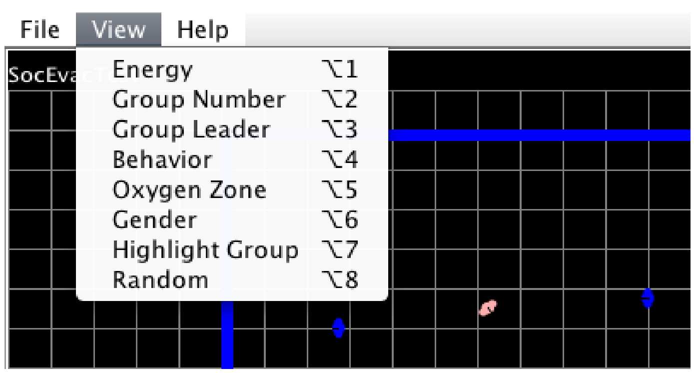

Figure 3.6: View Toggle

The view toggle shown above in figure 3.6 allows model users to select which attribute is used to display the color of each avatar in the model. Here, users can select from energy level, group number, leadership status, behavior type, oxygen available, gender, and can also choose to highlight one individual group.

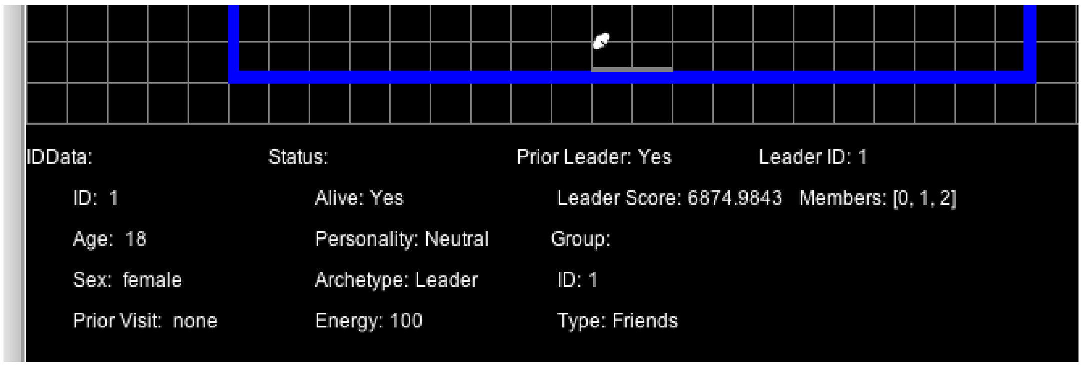

Figure 3.7: Real-Time Data

Figure 3.7 shows the real-time data that can be displayed in the second-generation model. This data is available for the simulation as a whole, or can be displayed for an individual avatar by clicking on that avatar during a simulation. The real-time data that is generated during a simulation by clicking on an avatar.

3.7 Summary

The first-generation model, originally intended to satisfy requirements for a social science theory based model, ultimately created more questions than answers, and the desire to address those questions led to the development and the specifics of the second-generation model. These differences led to the creation of two very different platforms designed to answer some of the same questions about the place of social science theory in modeling. Because these models are somewhat different in their own right, the similarities and differences between the two versions can be used to measure the effects of variable changes in each model.

The first-generation model stands on its own supporting the idea that social science theory has a place in building evacuation modeling, but the second-generation model is ultimately a better effort, and thankfully the engine for environment creation can be distributed in an attempt to inspire other model users to consider social science ideas in their modeling.

Chapter 4

SECOND-GENERATION SOCIAL BEHAVIOR MODEL

While the results from the first-generation model were promising, and show that adding group and crowd behavior elements may lead to a more accurate evacuation model of The Station, the platform raised questions about the appropriate amount of model variability, data validation methods, data privacy methods, display elements, and decision-making complexity.

Not content to settle for a single version of the modeling platform, the second-generation model improvements described here are an effort to expand upon the initial requirements for an agent-based model of evacuation of the Station nightclub accounting for social behavior and crowd density. This second-generation model addresses some issues with decision-making frameworks for avatars, shortcomings in visual display elements to make models accessible to all users, how results are created, stored, measured, and interpreted, and restructures some of the original hypotheses related to evacuation models.

4.1 Incorporating Economic Theory to Create Better Group Models

After development of the first-generation model and improved measuring and tracking metrics, analyses indicated new unresolved issues. The first attempt at creating model to account for social behavior was a probability-based system of agent evacuation.

While this system was fine for a conventional model, yielding overall results similar to the historical case study, results for individual agents were often extrapolations of the probabilities used. A rules-based system ran into similar problems. Results using flexible rules were not satisfactory at the individual level. With stringent rules, the results were similar to those of the probability-based system. Again, all of these results were accurate when total number of evacuees was measured yet they fell apart when tracking agents as unique individuals. However, by comparing model results for each agent and the exits and social pairing of real-life counterparts it is possible to validate models with more confidence.

In order to create a model with results that make sense on an individual level in a changing environment, developments in the economic field of game theory provide important insights. In some cases, strategies developed for games of incomplete information are an ideal fit. The strategies for games of incomplete information allow agents to prioritize decisions on individual and group levels, which are extremely important in this case (Bacharach 2006). This means that agents can act and interact in generalized ways because their instructions are not specific to any one situation.

There are several elements of economic theory that should prominent roles in future evacuation models. The first is the concept of the difference between absolute and relative utility (Ball and Chernova 2008). In many rule-based simulation systems, absolute utility is the only measurement regarded as important. The highest absolute utility for an individual is the “rational” choice. However, in social environments, higher relative utility may be preferred since it makes an actor better off relative to peers. In group social situations, relative utility is important enough that rule and probability systems based only on absolute utility are not adequate. The best result for the group may not be the best choice for the individual, and there are conflicts in individual choice. Research shows that situations involving family members and close friends can discount the value of individually optimal choices (Davis 2011). Models that do not account for group benefits cannot result in accurate portrayal of social evacuation scenarios. At the most detailed level of data analysis, the DRC used origin analysis and egress path analysis for individuals, and then compared individual ideals to group motivations and actual exit paths, a method supported in previous literature on evacuations from fires (Notarianni 2000). Game theory offers modelers a framework to examine evolving situations where participants may have asymmetric and incomplete information, conditions present in almost any building evacuation scenario (Jehiel and Koessler 2008). Economic theory allows multiple agents to account for the decision of others in changing environments, something that is extremely difficult to accomplish with probability or rule-based systems that only focus on individual behavior. This model makes absolute and relative individual and group utility matrices available to agents in situations where they would have information about other agents (such as groups working together in leader and follower relationships).

Developments in the understanding of bounded rationality play an important role in the proposed modeling platform. Models relying on probability or rule-based systems rely on agent "rationality", however in any dynamic environment perfect information is not available to all participants (Simon 1955). By applying modifications to the definition of “rational” choice, models can use economic theory to give avatars a smaller and improved decision set to make choices based on information available in the past or in real time (Nielsen 2009). Bounded rationality allows the model to function without any future event information available to avatars. Probability and rule-based systems often benchmark to actual event results, which considerably reduces the usefulness of models as predictive tools (Aguirre, El-Tawil, et al. 2011). This is a failure of probability and rule-based models; the incorporation of game theory should result in a significant improvement of model results. The class of models in this dissertation avoids the use of benchmarking, which means that model results are no longer only a function of actual event numbers of previous model output. By only using information available to agents in the timeline of the simulation models can be predictive (McCauley Bush et al. 2009).

Programming an understanding that a sub-optimal decision for an individual may be better than an argument resulting in delayed action allows for more efficient group movement. For groups that are unlikely to disband, calculating benefits from being agreeable may lead to groups that are more efficient. Agents in a group can compare likely results of individual behavior for all members of their group and likely group results. For many tight social groups, a better group result may be more desirable than the best individual result, a significant departure from conventional building evacuation modeling.

Refusing to use information only knowable after the fact (such as what happened to a specific avatar in real life) limits the benefits of some developments of game theory to the model. For instance, many components of backwards induction reasoning cannot be included because evacuations are single occurrence events (Binmore 2007). It is also not possible to incorporate many forms of learning strategies because they rely on repeated games and player experience in an environment. While these and perhaps other limitations exist, they do not prevent economic theory from offering significant improvement in evacuation model development. Instead, by using results from the model, it is possible to analyze agent decision-making after the fact. Previous efforts solved many of the coding issues related to agent group communication, but they do not extend these developments to "emotion", or caring about other agents or avatars (Brogan and Hodgins 1997). Economic frameworks allow asymmetric information to exist in the model, meaning that some avatars may be better informed than others at any given time.

4.2 Improving Visual Display of Evacuation Models

Creating models with multiple levels of social relationships and information sharing requires improvements in model visualization. In order to create models with interdisciplinary utility it is vital to create a system where practitioners from multiple disciplines can understand what is occurring in the model. This means that a graphics-based real-time output had to be developed side-by-side data cataloging methods. The second-generation model has new visualization capabilities. In real time, it is difficult to convey large amounts of constantly changing information (Pham et al. 2010). While it is easy to catalog these data in data sets for viewing and analysis later, it is also useful to create a system where information is conveyed as events occur in the simulation. For those who are not comfortable with statistics, effective model visualization is an important tool to gain understanding of how the model is working and what needs to change to improve it. As models with ever-greater levels of complexity were developed it became clear that improved visualization tools were needed (Bonato and Nowakowski 2011).

The variable of color is used as the primary means of real-time visualization, since many geographic strategy games allow users to select color to isolate and display information about objects (Silva, Santos, and Madeira 2011). The model has a system where model users can toggle between different visual display modes that change the colors of each agent based on the attribute examined. Basing color selection on ideal hues and visual impact (Cooper and Kamei 2002, Drevermann and Travis 1998), this system allows users to control the displayed attributes, such as groups and gender, and to convey the information in a meaningful way. The model also logs data for later use and detailed analysis. The toggle system is based on the premise that focusing on one aspect of a dynamic environment allows a user to more clearly interpret diverse results (dos Santos and Brodlie 2004).

Based on a meta-analysis of human factors studies, a well-designed visualization system increases understanding of a system (Chen and Yu 2000). The hope is that removing the usual barriers between disciplines such as sociology and computer science will result in a system where everyone is accountable and everyone has a stake in a successful model.

While analysis often focuses on data collected after the model runs are complete, real-time visualization is valuable for accuracy checking, communication of concepts, and model presentation. Since this model is designed to track multiple types of social relationships, simultaneous visualization of all avatar attributes would be complex. Additionally, visualization systems make it possible to visually measure the incidence of trends such as supra force over single model runs, or even averaged groups of simulations (Lee and Shen 2009). There are numerous agent attributes that are valuable to visualize to aid discussion about the importance of social science theory in human simulation, so the toggle system is an ideal way to communicate a large amount of information without overwhelming model users.

Previous modeling efforts and commercial packages attempt to account for visualization in various ways. Some toggle between two and three-dimensional environments and viewpoints in an attempt to give viewers intimate vantage points while still allowing for a full-field view. Some provide "flyovers" where the reference position is constantly changing. Some attempt to use different visual indicators, such as size, color, and pattern to communicate multiple attribute values. Many models use a combination of visualization methods, and the results can be very complicated to watch.

These models simply have too much information to convey in a real-time visual manner. World strategy gaming provided an answer to the unique visualization problem. In many world strategy games, one or more players attempt to lead geographic areas, controlling resources, populations, and physical environments. These games are similar to enormously complicated models where players control avatars. Players have to control for physical, economic, and political changes, often in real time. Because there is too much information to visually convey on a map all at once, these games incorporate a toggle system, where attributes can be isolated or layered together, allowing for quick visual updates that are very clear. A great example of this system is the massively multiplayer real-time strategy game World of WarCraft, which uses a toggle system to display large amounts of information on maps for players (Entertainment 2011). Using successful strategy games and academic research as a guide, a two-dimensional visualization system focuses on one area at a time displaying results on a constant scale for clarity (Wann and MonWilliams 1996, Colet and Aaronson 1995). For visually complex future models (such as stadiums or ships), rooms may be isolated instead of floors to keep the information scale manageable by viewers (Guo 2009).

After careful analysis of the above visual methods, this model adopts a two-dimensional overhead view that does not move because The Station environment is simple enough to process visually all at once (in more complex environments one floor or area is highlighted at a time). The visual output was developed with a task-based system for visual clarity, and in the future, it may be possible to incorporate eye-tracking technology for two-dimensional spaces to measure user focus on the dynamic parts of the model (Tory and Moller 2004, Pineo and Ware 2012). Users can select from a variety of toggles that change the color of each avatar to convey relevant information. In order to provide maximum clarity and continuity, no other avatar attributes change. This allows users to quickly switch between toggles and easily keep track of individual agents. A pause and start function is also included, to allow users to stop the model to examine attributes and then continue from the paused point without interruption to the model. The following paragraphs illustrate some of these options.

The model will have several visualization toggles. The first is a display of Energy level as seen in Figure 4.1. In the model, each avatar has an energy level that changes based on smoke and fire damage. By selecting the “Energy” toggle, avatar colors are displayed from green (healthy) to yellow (injured) to red (seriously injured). If avatars become inactive, they are displayed as a gray color. This toggle will serve as a real-time indicator of the injury status of each avatar, and can assist model users in locating high and low risk areas and avatars. This conventional color system will allow users to quickly and clearly identify avatar status.

Figure 4.1: Energy Level

The next toggle is a Group Number identifier as seen in figure 4.2. Each avatar will adopt an identical color to their other group members. This toggle will show the number of separate social groups in the model, and allow for a quick visual indicator of the number of active social groups. This indicator is not as clear as the others are but helps differentiate the 144 unique social groups in the model.

Figure 4.2: Group Number Identifier

Figure 4.3 displays the Leader Behavior toggle, which highlights group leaders in white, and all other avatars in black. This allows a user to see which avatars are responsible for the exit paths of themselves or others, and which avatars are following the instructions of a group. This enables a user to identify if an avatar is charting its own path or following one chosen by another avatar.

Figure 4.3: Leader Behavior

The Behavior Model toggle in figure 4.4 is the default toggle at the initialization of the model. Each avatar is highlighted based on their behavior profile at the time. Avatars that are in a leadership role are displayed in green to yellow to red based on their level of haste in exiting the scenario. Calmer avatars are shades of green or yellow, and avatars in dire situations are shades of red. Follower avatars have a pink color. Avatars that do not have free will of movement due to environmental factors and crowd density (supra force) are colored light blue. This system allows a user to identify the behavior profile of any avatar; it highlights the avatars that are subjected to the individual, group, and supra-crowd force behavior models at any given time in a simulation.

Figure 4.4: Behavior Model

The Oxygen Zone toggle highlights avatars in areas of increased oxygen in blue, and all other avatars in gray. This provides a very quick indicator of which avatars are in zones of increased survivability, explaining decreased injury and longer active periods.

Figure 4.5: Oxygen Zones

The Gender toggle highlights female avatars in pink and male avatars in blue. This allows a user to see the geographic segregation by gender, particularly at the beginning of The Station scenario, where a higher proportion of males were in the bar area and a higher proportion of females were near the stage, perhaps the area of greatest danger in the Station scenario. This toggle allows a user to confirm gender composition of the modeled gathering. According to Henderson, gender composition is important because previous research indicates there is, “inhomogeneity in crowds due to sexual differences” when high percentages of a crowd correspond to one gender (Henderson 1971).

Figure 4.6: Gender

The Highlight Group toggle highlights the members of one individual group in white, and all other avatars in black. Highlight Group will single out one new group each timestep; it changes constantly. Here, there is a two person group highlighted attempting to exit via the main door. This allows a user to determine the geographic spread of group members, and serves as an uncluttered visual indicator to show group members converging over time.

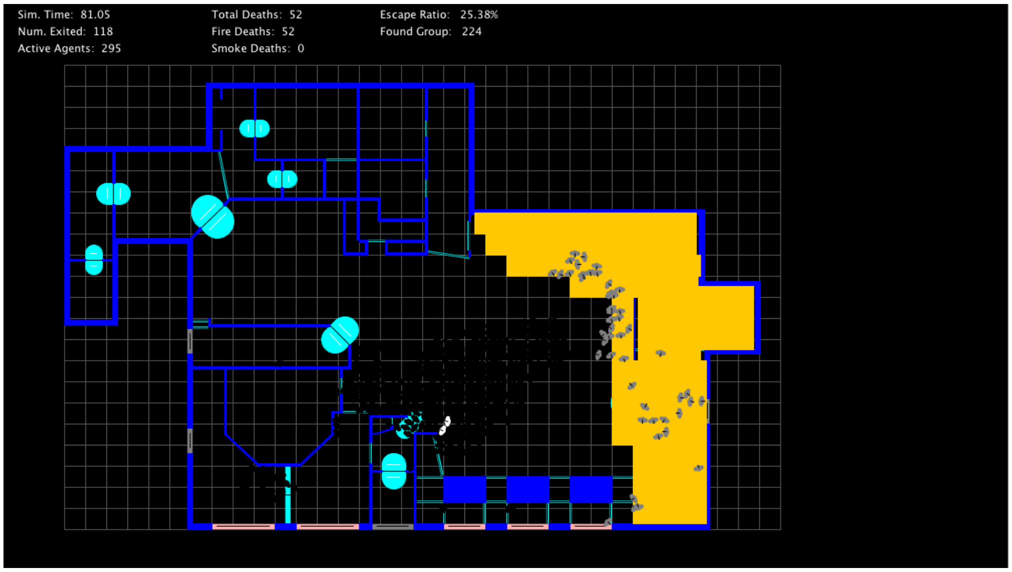

Figure 4.7: Highlight Group

The final two buttons in the model are Pause and Start, which will stop and restart the simulation so users can explore the model visually at their own pace. This allows for detailed inspection of the model. The model can be paused, making it possible to identify specific avatars, and then users can move through the various toggles to identify all of the relevant features of individuals or groups. By pausing the simulation, analysis can be accomplished in a manner that is not rushed, and then the model can return to running in real time when information is identified.

The second-generation model also builds upon previous research and incorporate a grid system for agent visualization. In the first-generation model, agents can see the environment unless there is an obstruction such as a wall or other avatars. Using a grid, it is possible to selectively limit visibility depending on light available, smoke incidence, and other factors in the areas of the Station. This will allow avatars to ignore unlikely hazards and opportunities (Chu et al. 2010, Song et al. 2006, Yang et al. 2004). However, reduced visibility may also change agent behavior if avatars are unable to see their group counterparts.

4.3 Model Accuracy:

The individual components put together in the first-generation model work together to create a dynamic model where the environment, individual traits, group traits, and density all affect evacuation choices. Put together, these different considerations result in a class of models that is much more lifelike than a simple closest-exit algorithm. In this model, avatars consider their group members, various hazards, population densities, and improvised exits when choosing the optimal path of exit. The optimal choices change constantly based on the specific conditions in the model, allowing avatars to make choices that are more informed based on new information. By comparing behaviors and dynamically approaching the environment, agents in this model are able to deal with uncertainty by constantly reevaluating choices.

This class of model should be accurate throughout simulations, not just at the end, and the additive process strives to create models that make sense at any time during the simulation. The included traits and resulting situations allow avatars to explore multiple possible exits, successfully remove themselves from hazard areas, and negotiate environments together with their social peers. Testing for this desired level of accuracy is one of the social science-based original contributions of this work.

This multiple-layer model is unique because of the simultaneous addition of hazards, environment, individual traits, group-level traits, and supra-crowd force. With these conditions included, the model produces results that are significantly different from versions that only include individual behavior. These traits are important for model builders to consider and include in future models, because human beings do not operate in social vacuums.

4.4 Measurement and Hypotheses:

Because most of the patrons in The Station were members of social groups, the group component is paramount in a model of the evacuation. Aguirre et al. 2005 found that most of the victims of the fire considered their social bonds initially more important than evacuation, and instead of immediately leaving the building, they searched for loved ones and friends, sacrificing precious seconds (Aguirre 2005). A thorough analysis of The Station fire indicates that without accounting for social behavior it would not be possible to replicate results in an accurate manner. This model makes it possible to gauge how often avatars are acting in their own interest or that of their groups.

Most evacuation models focus only on the numbers of overall evacuees and dead victims. If the goal is to understand what happened and why, this is simply an inadequate benchmark. It is very easy to create a model with an appropriate survival percentage, but that model will not necessarily be accurate in any other way. This dissertation offers a more detailed method of model measurement, and the creation of "point accurate" models that have the potential to conform to their scenarios at any point, not just at the end of a simulation. Instead of measuring the overall survival percentage to determine model validity, the model can track every specific agent individually, and benchmark results to detailed actual occurrences. This results in very detailed model benchmarks instead of a single survival percentage statistic, and allows model users to have a much better insight into the inner workings of the model.

This type of detailed analysis allows users to have a much better idea of how accurate a model is, how diverse the results are, and how parameter changes influence model functioning and results. This tracking results in a much more complete model dataset that users can analyze (and also creates new questions about validation to be discussed in chapters 4 and 6). This detailed data also allows for a much more complete model comparison, and makes it possible to test several hypotheses that could not be confidently analyzed with only survival percentage tracking.

The deliverables of this dissertation led to four hypotheses about the differences between conventional evacuation models and my model in scenarios with high group association such as the Station.

One hypothesis of the group behavior model is that avatars who are members of social groups will take a longer time to exit, and therefore be more likely to perish in the simulations where a quick evacuation is necessary to survival. In The Station scenario, a quick evacuation was required to exit the building without significant injury or death. The environment became inhospitable to human life within four and a half minutes, making any deviation from an exit hazardous to the health of the victims (Gill et al. 2010).

A second hypothesis is that the group behavior model can better explain the specific exit paths that victims used. Evacuation models using closest exit or fastest exit guidelines, including prior models of The Station, do not do a good job of estimating the exit that victims actually used in the fire. Because many victims initially moved towards other group members and often away from the nearest available exit, closest-exit models are inadequate. A model accounting for these phenomena should do a better job at estimating the actual results of an evacuation where social behavior was an important factor.

A third hypothesis is that by using physical parameters that do not account for human intelligence, conventional evacuation models have results that are too homogenous between runs, and will show decreased accuracy during model runs. If distance, density, and flow rates are the primary evacuation decision tools of agents, each agent is likely to take the same exit path most of the time, resulting in very regimented models that do not actually account for the situation that took place. A model accounting for more advanced behavior should yield a wider range of results. This hypothesis is tested by comparing exit path diversity for each agent with group social behavior turned on and off.

A final hypothesis is that a model accounting for social behavior elements will have better fitting results when tracking agents to final exits. Agents in a model accounting for social behavior should be more likely to reproduce their real-world counterparts' actual exit choices. Similarly, avatars that died in the simulation should more closely match the people that perished in the Station fire.

The second-generation ABM developed for this dissertation makes it possible to test these hypotheses, as discussed in chapter 5.

4.5 Model Validation

An important part of understanding the contributions of this study is whether or not it is possible to gauge the effect of adding group social behavior and supra-crowd force. Validation ideas and the improved visualization system help model users explore and compare model runs in detail. Despite the potential value of ABM simulation methodology for improving building safety--and more generally for developing new scientific knowledge--there has been a general lack of validation of ABM models of building evacuations as well as a tendency to adopt irrational, panic explanations of human behavior. The decision rules used to guide the action of avatars in other ABM simulations are not grounded in tested assumptions about the social behavior of people and their collectivities, their decisions to evacuate, and their evacuation behaviors in buildings on fire. It makes many ABM simulated outcomes unreliable as well as un-interpretable. Model validation is a difficult task because in most cases the level of detail does not match the information provided about the Station (Robinson and Brooks 2010). While there are established methods developed to track basic models, they generally do not address the prospect of tracking individual agent results or motivations (Valckenaers et al. 2007). There have been previous efforts to create certification systems or general guidelines of model validation (Youngblood and Pace 1995, Lord et al. 2005), but they have not been widely adopted by the modeling community.

Thoroughly validated models should provide accurate replication of historical processes and show that agents respond in a realistic manner to their circumstances. It is perhaps one of the most difficult dimensions to implement in ABMs of building evacuation. This dissertation includes a description of model validation attempts in chapter 5, and suggestions for the future in chapter 6. Because other models of The Station do not account for individual outcomes, it would not make sense to benchmark my results to prior work (although this is done to show the difference social science theory-based modeling creates in results). This validation includes novel methods of analysis used to determine accuracy of agent paths, relationships, and motivations behind agent actions. This allows for tests of social science theories in addition to survival percentages.

4.6 Leadership

It may also be possible to improve the study of leadership. Leadership components have been included in some models, but it may be possible to complement previous efforts by introducing dynamic leaders (Luo et al. 2008, Pelechano et al. 2005).

In an attempt to create a model with dynamic leadership functions, a leadership potential equation with values for group role, gender, age, building familiarity, proximity to exit, proximity to other group members, and prior leadership was created. It complements the static leadership model created as a part of the “Interaction Between Building and Occupant Responses During Collapse” (IBORC) project, funded by the National Science Foundation.

By comparing otherwise identical versions of the model with and without social behavior elements, it is possible to determine how leadership and group membership influence model results. Previous study showed improvement in evacuation time with “trained” leaders and communication among agents, but this concept needs to be coupled with changing leadership and social relationships that result in bonds between agents (Pelechano and Badler 2006).

Distance from other group members, distance from exit points, and injury are all dynamic values that change with the environment. Each group member calculates the square of distance to each other group members. The lower the sum of the squared values (least squares) the closer and more central each group member is to the rest of the group. A lower value increases the likelihood that a group member will be a group leader. Group members that are very far away from all other group members will be more likely to leave the group. Distance from the closest exit point also increases the likelihood of group leadership, because the avatar would be close to an actionable exit. Injury decreases the likelihood of group leadership, because an incapacitated group member would make for a suboptimal group leader.

The equation to calculate group leaders in the model is shown below in the Java language. A prior visit gives an agent a relatively large score premium. Smaller distances to other group members and smaller distances to exits are both positive leadership attributes. More “energy” or “health” is a positive value. Age is most desirable around 30 based on the demographics in our Station scenario (this value would very likely change in any other scenario). The final value is the premium from being a leader in the last timestep, to prevent constant leadership change.

```
leaderScore =
        2500* priorVisit
        + 10000*(1/(1+Math.pow(minDistance, 1.5)))
        + 10000*(1/(1+Math.pow(minExitDistance, 2.0)))
        + 3*(Math.pow(status.energy,1.2))
        + 1000 * idData.sex
        + 20 * (30 - Math.abs(30 - idData.age))
        + .1 * leaderScore * chosen;
    return leaderScore;
    }
```

The idea is that leadership should not necessarily be a variable defined before a model run, but it should also not change constantly. Usually, once someone is selected as a leader, things need to change significantly before another leader is chosen, something like the incumbent effect.

Once a leader is selected, other members in the group will target the group leader if they cannot see the leader’s exit goal, or the exit if they can see it. Willingness to break away from the group depends on all of the same factors as the leader calculation. For instance, an agent in bad shape right next to another exit will not ignore a sure thing to go to the other side of the building with a group.

These factors result in a modified leadership equation where the group leader can change dynamically even if the group members do not change. This is a significant improvement from the original model and better approximates reality.

These additions to an agent-based model of building evacuation are intended to result in a model that has more flexibility to attempt to explain complex real-world evacuations. Intended to be an interdisciplinary platform, the changes recommended here are designed to improve both accessibility for model building and analysis ability. While many of these features are exclusive to the SocEvac platform, it would be easy to add them to other popular modeling platforms.

Chapter 5

RESULTS

The second-generation model was developed for the task of modeling the Station nightclub as a single event. However, it is important to analyze the results in a manner that can provide insight to other modeling scenarios in the future. In this chapter, results are discussed for the final version of the model of the Station presented in this dissertation. These results are then compared to the actual event numbers to determine accuracy, and compared to other models of the Station (including my earlier efforts) as an academic exercise to facilitate a discussion about what modifications might benefit other modeling platforms in the future.

The results are presented in the Level One, Two, and Three formats. The Level One calculation only measures the total number of survivors to allow for a quick view of the model results. The Level Two analysis tracks overall use of every exit option, making it a little more useful for building and relationship analysis. Level Two is generally the output of most commercial modeling platforms. Level Three tracks every avatar as an individual, overly detailed for many model builders, but vitally important to determine if the social behavior functions in this dissertation are working correctly.

This chapter goes even deeper into the model output, tracking individual "fates" over time. While this level of analysis is not necessary to determine whether or not the model objective is working, it is pertinent to examine all of the possible output data to try and learn as much as possible about the benefits of adding social behavior code to models.

This is a study of what is possible with building evacuation model data, and whether or not we should devote resources to analyzing models in more detail. The chapter concludes with a discussion of model validation and accuracy. At this level of detail, there is no longer one best answer about what makes a good model, and a discussion of diminishing returns is important. While adding social behavior is an important contribution to micro-level evacuation models, it is necessary to start a conversation about what constitutes a significant difference in model results.

5.1 Final Model Datasets

The final model dataset used for the data analysis in the three levels consists of 1,000 runs of the “medium damage” profile of the second-generation model of the Station. The medium damage profile represented the best possible guess about fire and smoke damage in the Station based on the available data (Gill et al. 2010). In many ways, this output did not return the best model results, but it represents the most accurate model input.

One run of the model was corrupted, returning no results, and it was removed from the sample. There was one other run with a high death count that made it a clear outlier, however this run was included. To preserve the academic integrity of this first effort, this work includes all results that are the output of a complete modeling run. While outliers should generally be included in samples of complex modeling runs, output from corrupt code should be removed from samples.

Later in this chapter, two other 1,000 run datasets are included, “low damage” and “high damage” profiles. These are examples of results if the damage profile is modified, included as tests to show that incremental changes in input lead to incremental changes in output. Over the course of my dissertation research, tens of thousands of model runs of the Station were used to confirm the changes in model. Unfortunately, many of these early efforts were not standardized, were not recorded correctly, or were not final versions. As a result, only final datasets are used for comparison purposes to make sure that every parameter that returned a given set of results is accounted for correctly.

5.2 Level One Analysis

Level One represents the most rudimentary calculation, the total number of survivors. While this does not require a long discussion, it is worth covering because while this number might not be the most important benchmark in detailed micro-level modeling, it is often the only output considered in macro-output platforms. The overall survival level is the simplest possible benchmark to use for a hazard evacuation model. For evacuation models without a fatal-level hazard, clearance time is probably the best Level One measurement.

5.2.1 Purpose of Level One Data

Level One data is represented as two numbers, surviving avatars and perished avatars, or one survival percentage. In this model, Level One data serves as an executive summary, giving someone interested in the model a quick confirmation that the model is running. Oftentimes, the only reported model data consists of survival percentage and clearance time, making Level One data sufficient for this task. Level One data also has the side benefit of allowing these results to be compared to other modeling platforms that do not track any detailed avatar information.

5.2.2 Detailed Results

The Level One data is reported in the model as the number of avatars that perished in the Station model. In the 999 run set, an average of 131.56 avatars perished, with a standard deviation of 36.00. The median value was 129 deaths, and the mode value was 108 (which occurred 16 times out of 999). The dataset ranged from a low of 62 deaths to a high of 266. The second lowest value was 63 deaths, and the second highest value was 230. The inter quartile range of the dataset was between 103 and 155 deaths. Although the single 266 death outcome is an outlier, it is included because it was the result of a correctly functioning run of the model. The skewness of the death range was 0.454, indicating a positive or right-tailed skew of the distribution.

5.2.3 Comparison of Model to Actual Station Event

Compared to the actual Station event, the Level One results of the model are close but far from perfect. In the actual event, 100 participants perished in the fire at the Station. While SocEvac returns more deaths than actually occurred, the actual death count in the Station was within one standard deviation of the mean in 1,000 simulations. The data are skewed making some basic tests difficult to confidently discuss, but the takeaway is that the actual result is a plausible outcome of the model. 22.2% of the model runs returned results with less than 100 deaths, making the actual death numbers fall just slightly out of the inter-quartile range.

5.2.4 Comparison of Social Behavior Model to Closest Exit Calculation

The DRC decided to plot the closest exit for every avatar and use this as a basis to compare to my results, since the reason for originally creating a social behavior model was to offer an alternative to closest exit modeling. Using a strict closest exit calculation, 50 avatars would have perished in the fire and zero in the smoke according to my calculations. This is because no avatars would be moving in the opposite direction of each other, eliminating the biggest cause of congestion in my model. A simple comparison of SocEvac versus a pure closest evacuation calculation shows that SocEvac returns too high of a death rate, although the actual results are a reasonable outcome of the model. Meanwhile, the simple closest exit calculation returns a death rate that is too low.

A Level One analysis for a perfect model would fall somewhere in the middle of these two choices. While the 100 deaths for the Station scenario was completely plausible, it is reasonable to assume that result would have been on the higher side of the normal distribution of possible results. SocEvac returns a distribution where 100 deaths for the scenario would be in the lowest quartile of results, indicating that the design of SocEvac may be too "deadly". This makes intuitive sense when a user understands the limitations of SocEvac versus behavior of the victims in the incident. With the same hazards and environment, more intelligent avatars would likely be able to evacuate more efficiently that the avatars created in SocEvac.

5.2.5 Comparison to Other Models of the Station Event

There are a number of models that output Level One data for the Station nightclub fire scenario. These include Spearpoint's (2012) network model, Ewer et. al.'s (2008) evacuation model of the Station, and Pan et. al.'s (2006) MassEGRESS platform. While many of these models do not include the fire hazard directly, it is still worth comparing output of different model designs.

Spearpoint's effort, the latest in a series of pulished Station models, uses a networking model to create an analog of the Station nightclub (Spearpoint 2012). This model allows for only four exits (the doors at the Station) and assumes that those left in the building after 90 seconds will succumb to the fire. Spearpoint's model shows that 283 avatars are in the building at 90 seconds, and then subtracts 101 from these results (note that Fahy et. al. 2011 determined that 101 avatars used the windows as exits). This leaves a prediction of about 182 deaths from the fire. Assuming a fifteen second delay (proposed by Ewer et. al. 2008), the network model predicts 145 deaths.

While the Spearpoint effort is impressive, and network modeling is a promising improvement to complex building evacuation models, there were many assumptions made about the Station in this model that could be added to the simulations. According to our data, it was possible to survive longer than 90 seconds depending on the location of each avatar, and the results from the windows could be modeled, not assumed. Some versions of the model also preordained exit choices for certain avatars, which is an inaccurate way to try and model complex evacuations.

Spearpoint offers two output scenarios, based on the methodology used by Ewer et. a. described below. Spearpoint's unedited model returns 180 fatalities, and a model with 15 second fire delay returns 145 fatalities.

Ewer et. al. used the buildingEXODUS model and the SMARTFIRE model to create model of the Station evacuation (Ewer et al. 2008). Ewer et. al. offer three scenarios to the Station evacuation, resulting in 0, 180, and 84 deaths. The two scenarios examined here account for a fire starting instantly, and a fire starting with a 15 second delay. These results are much different than the delay added in the Spearpoint model. It shows that a relatively small change in the model can lead to a large change in results, which I believe is a possible indicator of over fitting in a model which will be discussed in the validation section. While Ewer et. al.'s effort is an excellent technical model, I think it would benefit greatly from the addition of complex social behavior elements. Ideally, codes like those produced in SocEvac would be merged with mature fire analysis platforms like SMARTFIRE, though the modeling of fire propogation is outside of the scope of this dissertation.

Pan et. al.'s MassEGRESS platform returns very different results, based on different assumptions about occupants and the building. The model returns only 89 deaths, making it one of the least lethal simulations of the Station (Pan et al 2006).

Three sets of preliminary results of the first-generation model developed for the Station scenario at the DRC are added to the comparison. As shown in the table below, two of the three Level One results from these early models are actually the most accurate using an overall survival rate analysis.

Table 5.1 shows a Level One analysis of the Station nightclub fire for the actual event, a closest exit model, the Spearpoint Network Model, The Ewer group model, MassEGRESS, and SocEvac. All values are mean values and ranges are not included because they are not always reported. SocEvac estimates a mean of 132 deaths, 32% higher than the actual 100 deaths in the fire. While SocEvac overestimates the overall death count, it appears to be within the acceptable range of outputs compared to peer models based on a Level One analysis.

Table 5.1: Level One Output

Model	Deaths	Error % (ABS)
Actual Event	100	
Closest Exit	50	50%
Spearpoint Network	182	82%
Spearpoint Network Delay	145	45%
Ewer	180	80%
Ewer Delay	84	16%
Pan MassEGRESS	89	11%
First-Generation No Groups	108	8%
First-Generation Weak Groups	104	4%
First-Generation Strong Groups	137	37%
SocEvac	132	32%

5.3 Level Two Analysis

Level Two represents the current state-of-the-art in many models, tracking the number of avatars that use each exit in the model. This level presents similar results to other current commercial modeling efforts that track pedestrian flow levels and are used to make suggestions about improving building design and evacuation plans. The Level Two analysis makes it possible to suggest building design changes when examining models in a predictive capacity.

5.3.1 Purpose of Level Two Data

In the Station nightclub, there were seven possible outcomes for each avatar opposed to the two options in Level One. There were four doors and two sets of windows creating six exit options, while the seventh option is succumbing to the fire during the simulation. SocEvac combined the five windows into two groups because individual windows were not identified in documentation of the evacuation.

The purpose of Level Two data is to evaluate which exits were used in the model, allowing model users to gather more information than is possible from a simple survival rate. A Level Two analysis allows analysis of model exits, comparisons of starting locations and exit preferences. It allows researchers to ask more complex questions about how models are working and why certain results are returned.

5.3.2 Detailed Results

In addition to the results for the number of dead explored in the 5.2 subsection, there were six other exit options analyzed for Level Two. This level of analysis looks at the aggregate exit results for the sets of windows in the greenhouse and bar area windows, and the bar, kitchen, main, and stage exits. Since we are now looking at more detailed model output, this analysis will go more in depth than the statistics looked at in Level One. When there is only a binary outcome, exit or not, it is not possible to learn as much about how a model is running. In contrast, when examining use of each exit, it becomes possible to see if exit relationships or preferences exist.

This analysis begins by examining the frequency of the use of different doors. The first is the Bar Exit. In SocEvac, the bar exit was used by an average of 90 avatars per simulation, with a standard deviation of 10 avatars. The minimum number that used the exit was 47, and the maximum was 118, with an IQR of 14 (83 to 97). There are four observations on the left tail that are statistical outliers. The skewness value is negative, indicating a longer left tail. The kurtosis value of 3.17 indicates that the tails are underweighted.

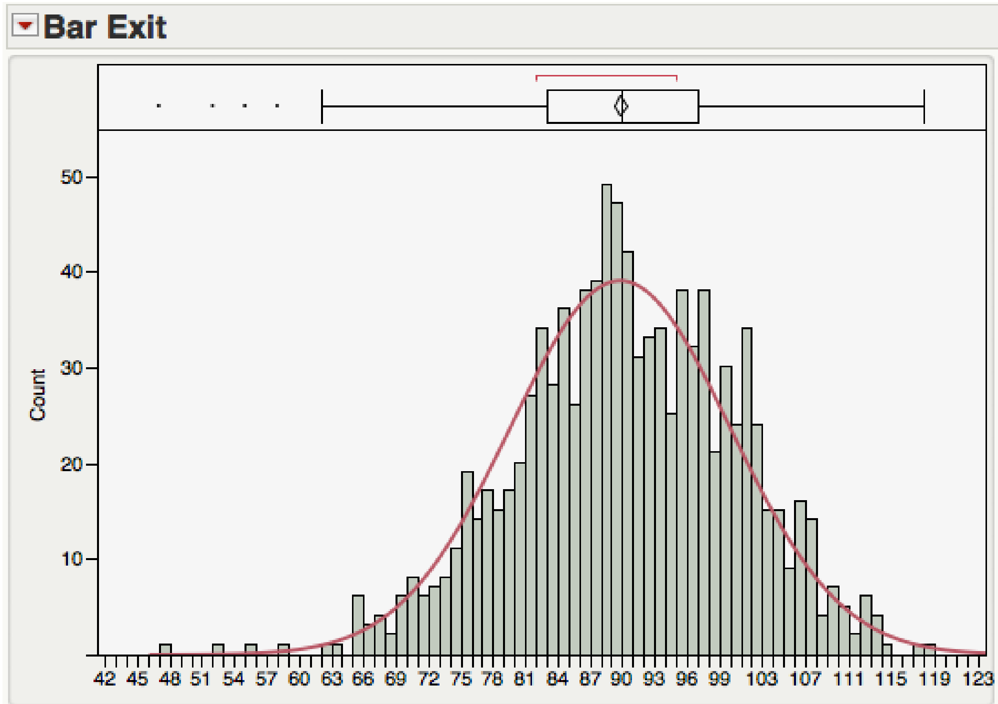

Figure 5.1: Bar Exit Distribution

As shown in the histogram above and supported by the descriptive statistics, the distribution of the use of the bar exit shows a relatively predictable distribution. The best distribution fit for the bar exit is the normal distribution. It appears that SocEvac functioned correctly when looking at the results of this exit.

The Kitchen exit appears to be an example of SocEvac functioning incorrectly. On average, only two avatars used the Kitchen exit to evacuate, with a standard deviation of one. The full range of avatars using the exit was between 1 and 9.

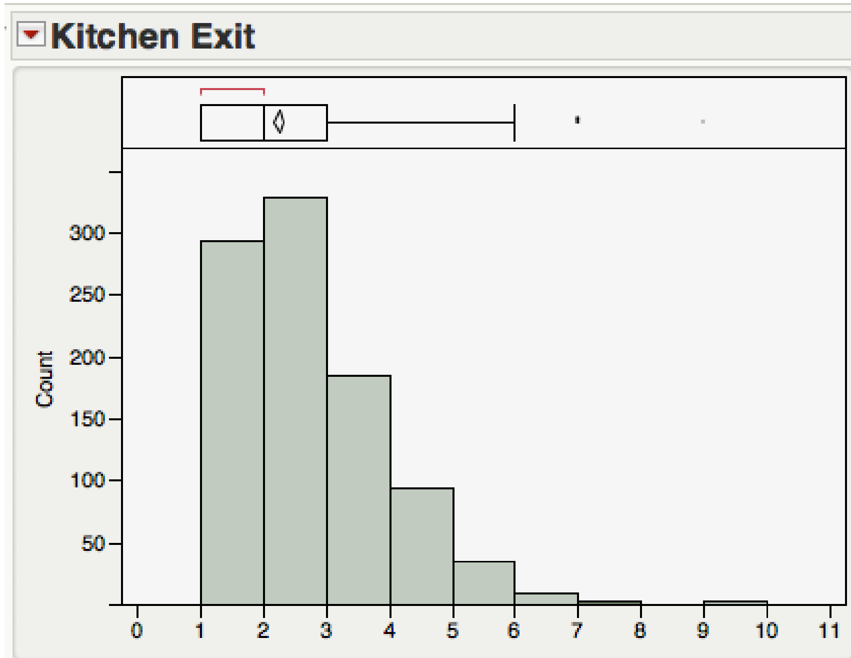

Figure 5.2: Kitchen Exit Distribution

The distribution of the Kitchen exit was not normal, and there is so little diversity of results that it is difficult to interpret much about the Kitchen exit. The best fit distribution for the kitchen exit is the binomial distribution, confirming that the normal distribution is a bad fit for this exit. In SocEvac, the Kitchen exit was only known about by avatars that had visited the Station before the evening of the fire, explaining the very limited use of the exit.

The Main exit was used by an average of 107 avatars per simulation, with a standard deviation of 17 avatars. The minimum number that used the exit was 28, and the maximum was 156, with an IQR of 19 (98 to 117). There are numerous observations on both tails that are statistical outliers. The skewness value is negative, indicating a longer left tail. The kurtosis value of 5.42 indicates that the tails are underweighted.

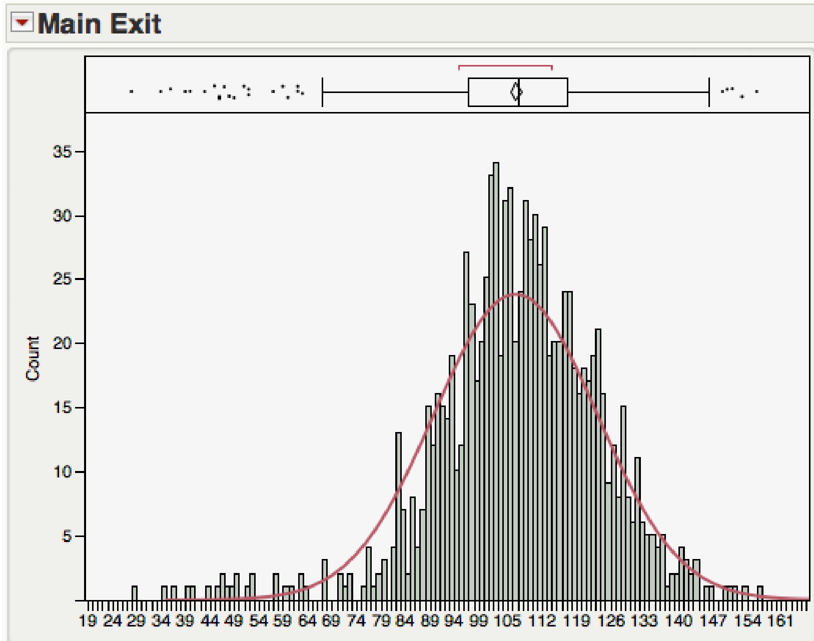

Figure 5.3: Main Exit Distribution

As shown in the histogram above and supported by the descriptive statistics, the distribution of the use of the bar exit shows a relatively predictable distribution, although the distribution is tighter than a normal distribution. A normal distribution is still the best fit for these results. It appears that SocEvac functioned correctly when looking at the results of this exit.

The Stage exit was used by an average of 36 avatars per simulation, with a standard deviation of 8 avatars. The minimum number that used the exit was 8, and the maximum was 55, with an IQR of 10 (31 to 41). There are two observations on the left tail that are statistical outliers. The skewness value is very slightly negative, indicating a longer left tail. The kurtosis value of 2.76 indicates that the tails are underweighted.

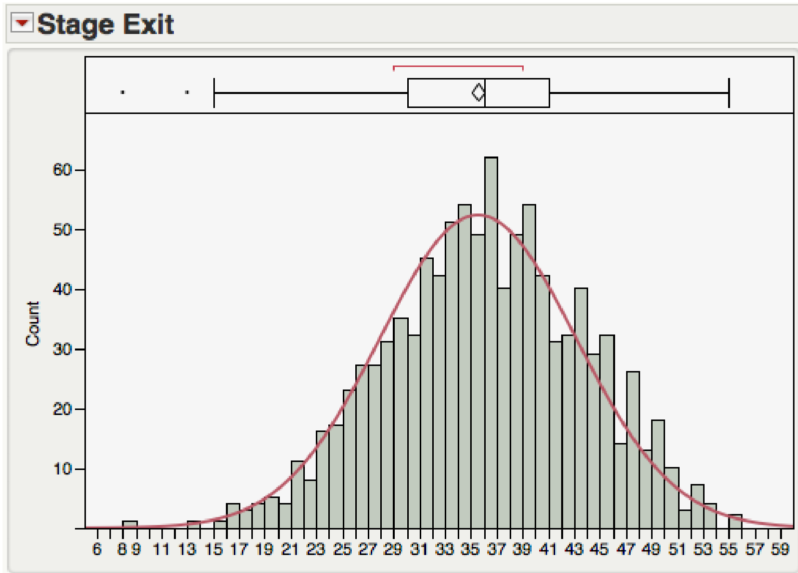

Figure 5.4: Stage Exit Distribution

As shown in the histogram above and supported by the descriptive statistics, the distribution of the use of the stage exit shows a relatively predictable distribution, although the distribution does not appear to be normal. The Poisson distribution is the best fit for these results. It appears that SocEvac functioned correctly when looking at the results of this exit.

The windows in SocEvac operate slightly differently than the doors. In the SocEvac model, a door does not change over the course of a simulation, but a window is an unattractive option to each avatar until one avatar breaches a window, at which point the exit increases in attractiveness. As a result, while it is reasonable to expect normal distributions at each door, the same may not be said for windows. The Greenhouse windows were used by an average of 26 avatars per simulation, with a standard deviation of 11 avatars. The minimum number that used the windows was 1, and the maximum was 56, with an IQR of 18 (17 to 35). There are no statistical outliers in the results. The skewness value is very slightly positive, indicating a longer right tail. The kurtosis value of 2.35 indicates that the tails are underweighted.

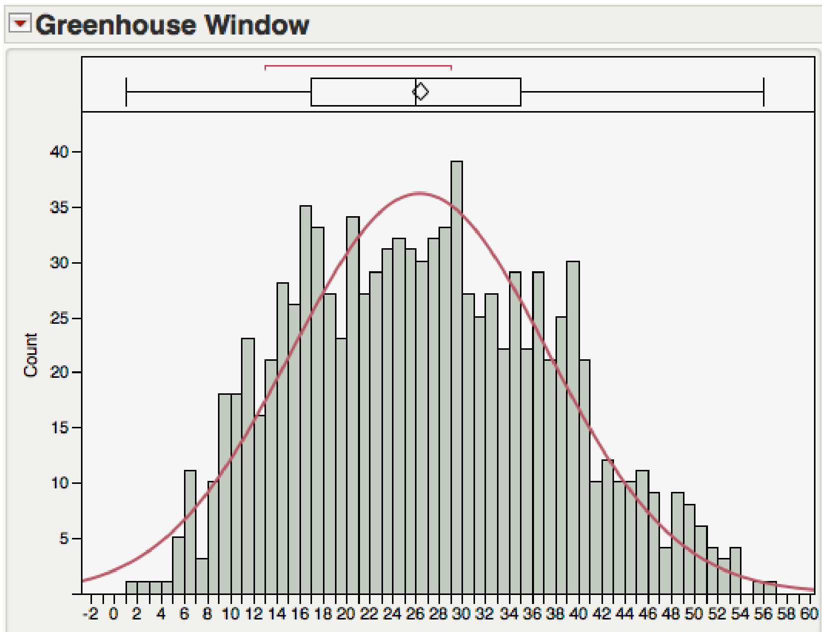

Figure 5.5: Greenhouse Window Distribution

As shown in the histogram above and supported by the descriptive statistics, the distribution of the use of the stage exit shows a somewhat predictable distribution, although the distribution does not appear to be normal. Interestingly, although the kurtosis value was positive, the histogram appears to show that the center of the distribution is underweighted with some overweighting on each side of the body of the distribution. Looking at the results of the two tails, it does not appear that these results are approximating a normal distribution. The log-normal distribution is the best fit for these results. It appears that SocEvac functioned correctly when looking at the results of this exit, accounting for the different functions of doors and windows.

The Bar windows were used by an average of 73 avatars per simulation, with a standard deviation of 24 avatars. The minimum number that used the windows was 12, and the maximum was 152, with an IQR of 34 (56 to 90). There are two statistical outliers on the right tail. The skewness value is very slightly positive, indicating a longer right tail. The kurtosis value of 2.64 indicates that the tails are underweighted.

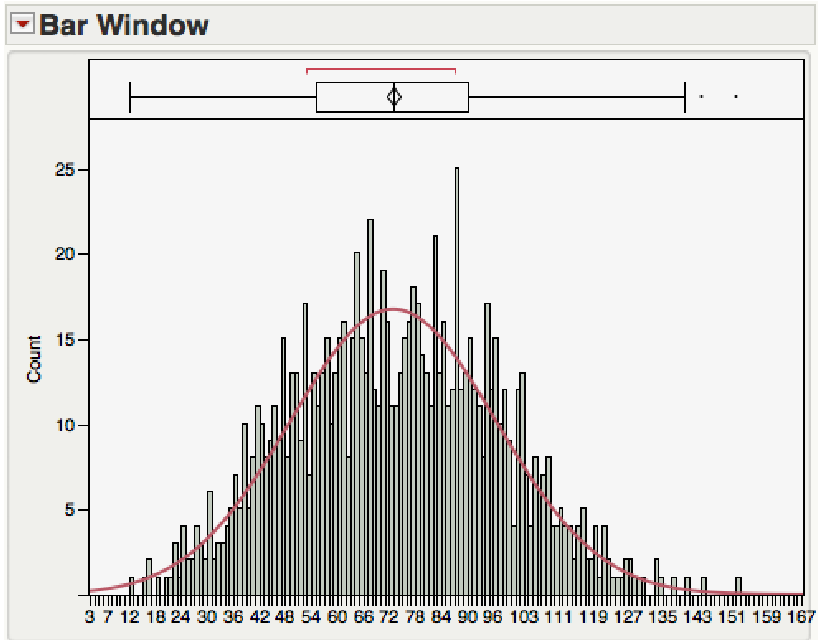

Figure 5.6: Bar Window Distribution

As shown in the histogram above and supported by the descriptive statistics, the distribution of the use of the stage exit shows a somewhat predictable distribution, although the distribution does not appear to be normal. Interestingly, although the kurtosis value was positive, the histogram again appears to show that the center of the distribution is underweighted with some overweighting on each side of the body of the distribution. Looking at the results of the two tails, it does not appear that these results are approximating a normal distribution, even though a normal distribution is the best fit for these results. It appears that SocEvac functioned correctly when looking at the results of this exit, accounting for the different functions of doors and windows.

As originally shown in section 5.2.2, the distribution of deaths was not normal. The results were positively skewed, with a kurtosis of 2.65, and the histogram does not look like it is approximating a normal distribution. The best fit for the deaths is the inverse Gaussian distribution, confirming the long right tail of the distribution.

According to two-sample T-tests (99% confidence), there are positive correlations for the following exit choices: Greenhouse window and Bar exit, Greenhouse window and bar window, Main exit and Kitchen exit, Bar exit and bar window, and Kitchen exit and bar window. While this might not make intuitive sense since exit choices are in essentially in competition with each other, it is clear that the death rate is not normally distributed, so at least in this instance of SocEvac, not using one exit does not necessarily mean an avatar is going to use another one.

There are four negative correlations (99% confidence): Greenhouse window and the Main exit, Greenhouse Window and death, Main Exit and death, and Bar Exit and death. Based on the specific model situation, and corroborated by films and other sources of information about this fire, there is a massive queuing at the entrance to main exit, so it makes sense that if more avatars are likely to use the greenhouse window, it would cause fewer avatars to either use the main exit or perish. The other two negative correlations are with door exits and death, making more sense than negative correlations between exits.

The first half of figure 5.7 shows an example of a 3-D bivariate histogram of the Main exit and Death, clearly showing the negative correlation. This is in contrast with the second figure, which shows the uncorrelated relationship between the Main exit and the Stage exit.

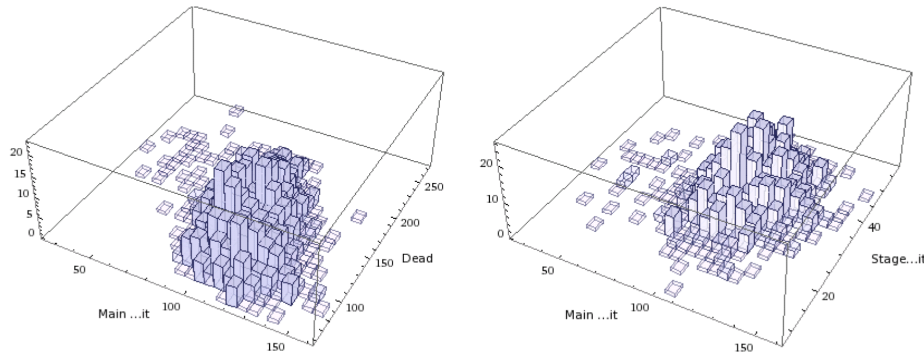

Figure 5.7: Bivariate Histograms

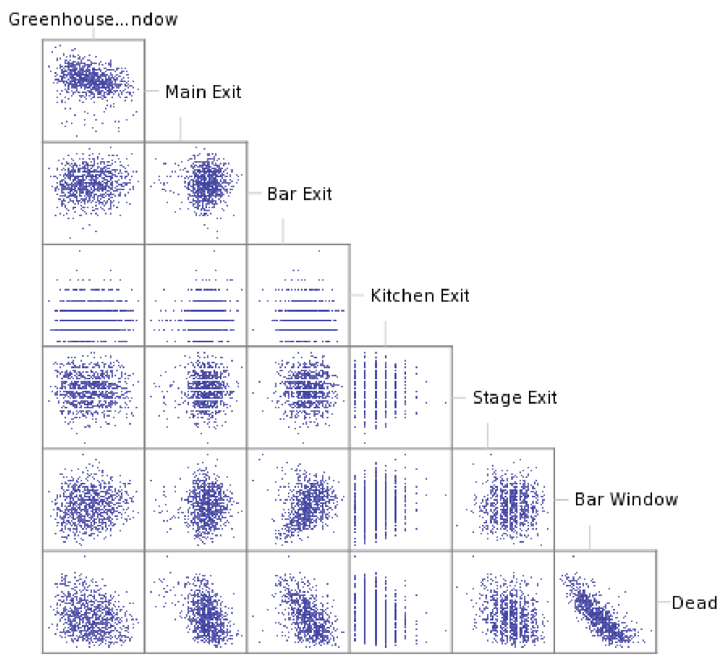

Figure 5.8: Exit Relationships

The scatterplot array shown above in figure 5.8 shows the two-dimensional relationships between each pair of results, confirming that there are not strong correlations between many of the pairs. The Level Three analysis in the next section makes it possible to explore the relationships between exits in the model of the Station in more detail.

5.3.3 Comparison of Model to Actual Station Event

The actual Station event had exit volumes that were similar to the mean output of SocEvac. In order of use, the six exits were unknown (13), Kitchen (17), Stage (24), Greenhouse window (34), Bar window (71), Bar (78), and the Main exit (128). The other 100 people in the building perished in the fire. Since there were 13 outcomes that are still unknown, it is not possible to have a "perfect" Level Two result, but it is possible to obtain results that are very close.

Based on the absolute value of differences between the SocEvac means and the actual event numbers, SocEvac returns a total error count of 101 Level Two differences. A perfect score would be 13. While the absolute numbers make it possible to quickly see the magnitude of errors, examining the standard deviation range of each exit distribution yields more interesting results. In this Level Two analysis, SocEvac exit and death results are all within on standard deviation of the actual results of the Station scenario, with the exception of the Kitchen exit (17 people using the Kitchen exit would be an extreme outlier in SocEvac). These results indicate that the model makes sense in terms of the actual results of the Station.

5.3.4 Comparison of Social Behavior Model to Closest Exit Calculation

The closest exit results for the Level Two analysis show a very different pattern of exit use in the Station scenario. The closest exit calculations underestimate the use of the Bar exit, Bar window, Death, Greenhouse window, and Main exit. It overestimates the use of the Kitchen exit and the Stage exit. In percentage terms, the two largest errors are the underuse of the Bar exit (11.5% of actual) and the overuse of the Stage exit (795.8% of actual).

In the Level 2 analysis it is rather clear that the SocEvac results are more like the actual results than the closest exit results. This is unsurprising since the original purpose of this work was to examine why closest-exit modeling could not adequately explain the results of the Station nightclub fire.

5.3.5 Comparison to Other Models of the Station Event

This subsection compares results from SocEvac and previous DRC efforts to analyses of the Station by MJ Spearpoint at the Department of Civil and Natural Resources Engineering at the University of Canterbury in New Zealand, NIST, Thunderhead Engineering, the University of Greenwich, and the Engineering Informatics Group at Stanford University (Galea et al 2008b, Grosshandler et al. 2005, Pan et al. 2006, Spearpoint 2012, Thunderhead Engineering 2012). These groups have all developed and compiled other models of the Station and published these results. There are many similarities between their work and the DRC analysis of the Station, as well as many profound differences. In cases where it was not possible to directly understand the methodology or results of these published models, the DRC modified the results in a manner that would make them more accurate. Any errors here are mine, and not those of the other model authors. All results are presented as mean values because ranges are not always included (and ranges are not possible in models with no variation).

Although perfect comparisons are not possible for reasons that will become clear in this subsection, there are a number of other model platforms that have performed various different Level Two calculations of the Station event. Unfortunately for this comparison focused on the Station nightclub, many of these platforms focus on clearance times and not on fatality levels, so results will not be perfect fits for the actual situation.

However, it is still possible to examine the proportional exit usage and the theory behind these models.

This section considers the actual historical results, my closest exit model, the three DRC first-generation models, Simulex, buildingEXODUS, MASSEgress, two Pathfinder models, three of the EvacuatioNZ efforts, and SocEvac. Other models have been used to simulate the Station scenario, however they did not output the results needed for a Level Two analysis, making a comparison here very difficult.

This subsection will include a comparison of all available platforms to both the actual model results and the closest exit results of the Station scenario that were developed based on the DRC data from the Station. In order to discuss both the accuracy of these models relative to the real scenario and the methodology used, the next paragraphs will compare the results to both the real outcome and the predicted outcome of the closest exit model.

Table 5.2 below shows the Level Two results output for the models of the Station compared in this subsection. There are major barriers to this comparison. The first is that there are four different numbers of agents in the different scenarios. Most of the other modeling programs based their numbers off of the 420 person occupancy limit of the Station, not the 465 agents that the DRC and Barylick found to be in attendance (of the other population values measured) (Spearpoint 2012, Torres 2010, Barylick 2012). Tubbs and Meacham estimated there were between 440 and 458 occupants in 2007, a range above the occupancy level but below the attendees identified by the DRC (Tubbs and Meacham 2007). Additionally, the results are compared to two sets of numbers, the 452 known agent outcomes, but there are 465 total agents. This makes it almost impossible to return perfect modeling results, but that is not a great concern since a perfect result with 465 agents and seven possible outcomes for each agent would almost certainly be evidence of overfitting. Models with an asterisk by their name have interpreted results based off of model output and author instructions.

Table 5.2: Level 2 Model Results

Model Results	Bar Exit	Bar Window	Dead	Greenhouse
Actual Event	78	71	100	34
Closest Exit	9	49	50	22
First-Generation No Groups	37	2	108	6
First-Generation Weak Groups	34	10	104	6
First-Generation Strong Groups	113	56	137	2
Simulex*	20	0	0	0
buildingEXODUS*	22	0	0	0
MASS Egress*	36	0	0	0
Pathfinder (SFPE)*	19	0	0	0
Pathfinder (Steering)*	19	0	0	0
Spearpoint Minimum Distance*	31	0	0	0
Spearpoint Assigned*	20	0	0	0
Spearpoint 90 Seconds*	22	70	145	31
SocEvac	90	73	131	26

Model Results	Kitchen Exit	Main Exit	Stage Exit	Total
Actual Event	17	128	24	452
Closest Exit	41	103	191	465
First-Generation No Groups	164	119	29	465
First-Generation Weak Groups	177	105	29	465
First-Generation Strong Groups	7	126	24	465
Simulex*	3	213	184	420
buildingEXODUS*	4	214	180	420
MASS Egress*	4	293	87	420
Pathfinder (SFPE)*	3	207	191	420
Pathfinder (Steering)*	3	201	197	420
Spearpoint Minimum Distance*	42	0	347	420
Spearpoint Assigned*	3	212	185	420
Spearpoint 90 Seconds*	17	129	41	455
SocEvac	2	107	36	465

Looking at the table, it is easy to see the large diversity in model results for the same scenario, even after accounting for the differences in fatalities and numbers of agents. Bar exit use ranges from 19 to 113, and numbers for Stage exit use range from 24 to 347. These profound differences indicate that the different model profiles significantly change results in the Station scenario. In some instances, such as the different calculations used in the EvacuatioNZ profiles this makes sense.

Actual event numbers and the closest exit model are detailed in previous subsections, and although the first-generation model results have already been discussed in chapter two, they are worth revisiting here. The first-generation model was almost certainly overfit. It was developed as a single-use model for the Station as a test bed for the social behavior model concept. The “no groups” and “weak groups” models returned very similar results, because we underestimated how strong group bonds would need to be before agents would choose to abandon groups in extreme situations. In the first-generation model, there is almost no discernable difference between “no groups” and “weak groups” because with several operations per second, even low probability events (1% per timestep [one twentieth of a second]) are likely to occur in seconds. The “strong group” version of the model returned much better results, signaling that the work that resulted in this dissertation was probably worth exploring.

Many of these results are very different from the actual results of the tragedy, primarily because these models focused on clearance time and not on the specific fatal hazard in the Station. As a result, many of these simulations have no fatalities. It is possible to add fatalities to these models based on active agents still in the building at time of certain lack of survivability, but these results are only modified in cases where the authors suggested how to do so (*).

Table 5.3: Model Differences from Actual Event

Absolute Difference from Actual	Bar Exit	Bar Window	Dead	Greenhouse
Closest Exit	69	22	50	12
First-Generation No Groups	41	69	8	28
First-Generation Weak Groups	44	61	4	28
First-Generation Strong Groups	35	15	37	32
Simulex*	58	71	100	34
buildingEXODUS*	56	71	100	34
MASS Egress*	42	71	100	34
Pathfinder (SFPE)*	59	71	100	34
Pathfinder (Steering)*	59	71	100	34
Spearpoint Minimum Distance*	47	71	100	34
Spearpoint Assigned*	58	71	100	34
Spearpoint 90 Seconds*	56	1	45	3
SocEvac	12	2	31	8

Absolute Difference from Actual	Kitchen Exit	Main Exit	Stage Exit	Total
Closest Exit	24	25	167	**369**
First-Generation No Groups	147	9	5	**307**
First-Generation Weak Groups	160	23	5	**325**
First-Generation Strong Groups	10	2	0	**131**
Simulex*	14	85	160	**522**
buildingEXODUS*	13	86	156	**516**
MASS Egress*	13	165	63	**488**
Pathfinder (SFPE)*	14	79	167	**524**
Pathfinder (Steering)*	14	73	173	**524**
Spearpoint Minimum Distance*	25	128	323	**728**
Spearpoint Assigned*	14	84	161	**522**
Spearpoint 90 Seconds*	0	1	17	**123**
SocEvac	15	21	12	**101**

Table 2.4: Square of Model Differences from Actual Event

Square of Difference from Actual	Bar Exit	Bar Window	Dead	Greenhouse
Closest Exit	4761	484	2500	144
First-Generation No Groups	1681	4761	64	784
First-Generation Weak Groups	1936	3721	16	784
First-Generation Strong Groups	1225	225	1369	1024
Simulex*	3364	5041	10000	1156
buildingEXODUS*	3136	5041	10000	1156
MASS Egress*	1764	5041	10000	1156
Pathfinder (SFPE)*	3481	5041	10000	1156
Pathfinder (Steering)*	3481	5041	10000	1156
Spearpoint Minimum Distance*	2209	5041	10000	1156
Spearpoint Assigned*	3364	5041	10000	1156
Spearpoint 90 Seconds*	3136	1	2025	9
SocEvac	144	4	961	64

Square of Difference from Actual	Kitchen Exit	Main Exit	Stage Exit	Total
Closest Exit	576	625	27889	**36979**
First-Generation No Groups	21609	81	25	**29005**
First-Generation Weak Groups	25600	529	25	**32611**
First-Generation Strong Groups	100	4	0	**3947**
Simulex*	196	7225	25600	**52582**
buildingEXODUS*	169	7396	24336	**51234**
MASS Egress*	169	27225	3969	**49324**
Pathfinder (SFPE)*	196	6241	27889	**54004**
Pathfinder (Steering)*	196	5329	29929	**55132**
Spearpoint Minimum Distance*	625	16384	104329	**139744**
Spearpoint Assigned*	196	7056	25921	**52734**
Spearpoint 90 Seconds*	0	1	289	**5461**
SocEvac	225	441	144	**1983**

Tables 5.3 and 5.4 above show the absolute difference in number of users for each exit between the output of each model and the exit totals gathered by the DRC. The second table shows the square of those values between each avatar fate in a model and the actual value discovered by the DRC. Analysis of absolute value differences treat every error as a constant, and the sum of squares of these values makes bigger differences exponentially larger than smaller ones. Looking at the first of the two tables, the closest end exit results using this measure are SocEvac, Spearpoint's EvacuatioNZ model (at 90 seconds), and the first-generation strong group model attempt. In this case, the Spearpoint model contains an assumption about the windows. Spearpoint's model uses an assumption from Fahy et. al. that 101 participants used the windows to evacuate. In this model comparison, 101 was subtracted from the total fatalities and assigned these avatars to the appropriate windows (Fahy, Proulx, and Flynn 2011). While this makes the results of Spearpoint's model much more accurate, this is appropriate in this case because the use of other exits is somewhat proportional to the real occurrence. Unfortunately, SocEvac and the first-generation models are the only platforms that account for windows at all, which makes it difficult to make some comparisons to other modeling efforts. At least Spearpoint considers the possibility of window use even though his model does not allow for them. When looking at the sum of squares of differences, the first-generation strong groups model pulls ahead of the Spearpoint 90 second model, because of the smaller differences. Interestingly, the Simulex, buildingEXODUS, MASSEgress, Pathfinder, and Spearpoint assigned models all return highly similar values. This suggests the possibility of benchmarking one model to another.

Table 5.5: Model Difference from Closest Exit Results

Absolute Difference from Closest Exit	Bar Exit	Bar Window	Dead	Greenhouse
First-Generation No Groups	28	47	42	16
First-Generation Weak Groups	25	39	46	16
First-Generation Strong Groups	34	7	13	20
Simulex*	11	49	50	22
buildingEXODUS*	13	49	50	22
MASS Egress*	27	49	50	22
Pathfinder (SFPE)*	10	49	50	22
Pathfinder (Steering)*	10	49	50	22
Spearpoint Minimum Distance*	22	49	50	22
Spearpoint Assigned*	11	49	50	22
Spearpoint 90 Seconds*	13	21	5	9
SocEvac	57	20	19	4

Absolute Difference from Closest Exit	Kitchen Exit	Main Exit	Stage Exit	Total
First-Generation No Groups	123	16	162	**434**
First-Generation Weak Groups	136	2	162	**426**
First-Generation Strong Groups	14	23	167	**278**
Simulex*	10	60	7	**209**
buildingEXODUS*	11	61	11	**217**
MASS Egress*	11	140	104	**403**
Pathfinder (SFPE)*	10	54	0	**195**
Pathfinder (Steering)*	10	48	6	**195**
Spearpoint Minimum Distance*	1	103	156	**403**
Spearpoint Assigned*	10	59	6	**207**
Spearpoint 90 Seconds*	24	24	150	**246**
SocEvac	9	4	155	**268**

Table 5.6: Square of Model Difference from Closest Exit Results

Square of Difference from Closest Exit	Bar Exit	Bar Window	Dead	Green house
First-Generation No Groups	784	2209	1764	256
First-Generation Weak Groups	625	1521	2116	256
First-Generation Strong Groups	1156	49	169	400
Simulex*	121	2401	2500	484
buildingEXODUS*	169	2401	2500	484
MASS Egress*	729	2401	2500	484
Pathfinder (SFPE)*	100	2401	2500	484
Pathfinder (Steering)*	100	2401	2500	484
Spearpoint Minimum Distance*	484	2401	2500	484
Spearpoint Assigned*	121	2401	2500	484
Spearpoint 90 Seconds*	169	441	25	81
SocEvac	3249	400	361	16

Square of Difference from Closest Exit	Kitchen Exit	Main Exit	Stage Exit	Total
First-Generation No Groups	15129	256	26244	**46642**
First-Generation Weak Groups	18496	4	26244	**49262**
First-Generation Strong Groups	196	529	27889	**30388**
Simulex*	100	3600	49	**9255**
buildingEXODUS*	121	3721	121	**9517**
MASS Egress*	121	19600	10816	**36651**
Pathfinder (SFPE)*	100	2916	0	**8501**
Pathfinder (Steering)*	100	2304	36	**7925**
Spearpoint Minimum Distance*	1	10609	24336	**40815**
Spearpoint Assigned*	100	3481	36	**9123**
Spearpoint 90 Seconds*	576	576	22500	**24368**
SocEvac	81	16	24025	**28148**

Tables 5.5 and 5.6 above show the same type of analysis for the differences (and the square of differences) between each model and my closest exit calculations (Aguirre et al 2011, Spearpoint 2012, Fahy et al 2011, Grosshandler er al 2005, Galea et

al 2008, Chaturvedi et al 2006). This is done to check how different many of the models are to a traditional static exit pattern. In many ways, this is actually a more interesting output than the comparison between the actual results and each model. In this case, the Pathfinder models returned the best results. Spearpoint's assigned exit model was close behind, however, Spearpoint's minimum distance model, which one would expect to offer similar results to my closest exit model, was one of the worst fits to my closest exit calculations. This suggests that there is significant variation in inputs in addition to variation in outputs of the different models of the Station nightclub fire.

Interestingly, Simulex, building EXODUS, Pathfinder, and the Spearpoint assigned models all have similar degrees of error compared with the closest exit model. This suggests that these platforms use a similar calculation to determine optimal exit (such as the popular A* algorithm).

SocEvac is not a particularly good fit in this case, which is promising. Returning results that more closely resemble the actual event than a closest exit calculation is a sign that the ideas implemented in SocEvac have a place in micro-level building evacuation modeling. These results suggest that many of the models of the Station evacuation may be slightly modified closest-exit designs, and that behaviors considered in the models do not overrule exit distance.

5.3.6 Level Two Criticisms

It is important to note that these models of the Station (including SocEvac) all make interpretations of the placement and number of participants that were in the building at the time of the fire. Based on extensive documentary research on this fire, the DRC estimated 465 victims, but it is worth noting that there is not complete agreement in the modeling community about this number. When developing SocEvac, the DRC performed many simulations changing avatar numbers (and therefore density), but there was little connection between these learning experiments and the attempted recreation of a past event. Similarly, the models differ in the way they treat windows. Evacuation through the windows was a major source of egress in the Station scenario, but most simulation platforms ignored them. In the DRC's count of Station evacuees, it was found that 105 occupants used one of the five windows to exit the building (Aguirre, Torres, et al. 2011). Ignoring a means of egress used by almost one quarter of participants is a major oversight. Most of the modeling platforms do not account for windows at all, and the Spearpoint model mentions that the method of accounting for windows is to subtract the actual number of window users from those still inside the building at 90 seconds (for the comparison above 101 evacuees were proportionally balanced between the two windows, without this balancing, the Level One and Two results are not nearly as close as the actual event numbers). In future modeling efforts, it is vital that model builders recreate environments that are accurate when they are attempting to recreate real-world events. All of the models of the Station nightclub fire have some combination of incorrectly placed agents, incorrect numbers of agents, incorrect exits, and no hazard (some of these errors are unavoidable). However, researchers need to make an effort to minimize these differences whenever possible, especially when "benchmarking" or comparing modeling platforms.

Most of the available models focus on clearance time, which can sometimes be misguided when attempting to understand evacuations. In a scenario like the Station, where survival in the building after about four minutes was deemed impossible in any location of the building, it makes little sense to me to analyze clearance times beyond a few minutes. The Station was a disaster because of the number of fatalities it caused, not because of slow clearance time (thus, a slower burning fire would not have resulted in as many people dying even if building clearance took the same amount of time). Many of these models could benefit greatly from including fatalities to give a more accurate picture of the results of a fire.

In evacuation scenarios, whether others are alive or dead can significantly affect victim behavior, so that completely ignoring the potentially fatal elements of a situation like the Station seems like a major modeling oversight. In an event like the Station fire, where a deadly hazard was seen by most participants as the main reason for the evacuation, it is important to include that hazard in a model. Based on the DRC analysis, the main barrier to exiting the Station was other people, and how much other individuals impacted the ability to exit the building, whether they were alive and well, injured, or dead. Not accounting for the hazard is a major oversight that makes detailed model results much less accurate than they could be.

Another problem with most of the models of the Station is overfitting. Modeling can never become predictive if researchers model past situations by using data that only could have been known after the fact in these models. Avatars in an evacuation model should not have access to data such as the exit they used in the real scenario of the well being of their relatives and friends in the gathering experiencing the crisis. In a well-fit model, what exit participants used in real life should have no bearing on avatar functions in an agent-based model. This is a critical issue that is overlooked to save time and make for better looking results. The decision to include data that only could have been known after the fact is one of the main roadblocks to the wider adoption of simulation. The eventual goal of simulation should be the ability to have models that are generalizable and predictive. Although we are far from that goal in social behavior models (or even complex evacuation models), it will be impossible for us to ever move the science forward if our inputs are based on the results of past events. In essence, this is an analog of “cheating” in modeling. A model that functions “perfectly” for only one scenario with no possible variability is useless as a predictive tool. In order for us to push the envelope with predictive models, we need to abandon absolutist instructions that prevent diversity in models.

The final problem, which is apparent after looking at the Level Two difference results is benchmarking. It is worrying that multiple different modeling platforms have seemingly identical Level Two results. In some cases, variability in modeling output is not even reported. For example, if a model of 400+ avatars with seven (or five) output possibilities does not have any variability, it is almost certain that it is overfit (especially with end result data discussed above) so that it it is too tightly benchmarked to prior model results of actual results. With such a complex scenario, significant variability would be expected. An incident like the Station should have sufficient complexity to prevent highly similar results from multiple modeling platforms. If all of these platforms were developed with different engines, one would expect to see more variability in the results.

If variability is an issue, the 1,000 run batches used to evaluate SocEvac are a good start, and it is frustrating that some model builders consider a single run of a simulation to be indicative of finished results. As an example, looking back at the tables above, the “best” Level Two run of SocEvac has an absolute value error total of 49 (13 would be a perfect replication based on our knowledge). On the opposite end, the “worst” run of SocEvac has an absolute value error total of 319. The minimum and maximum squared totals were 571 and 35259 (with the maximum making SocEvac’s average of

1983 look hopelessly optimistic). With ranges like this, even when extreme values are rare, it is irresponsible to only look at one result or at a small batch of results. Thankfully, some researchers such as Spearpoint report variability along with mean results, although unfortunately this not true of all platforms. Reporting variability of findings is a clear and simple way to advance simulation modeling.

5.3.7 Level Two Conclusion

The Level Two analysis is very telling; SocEvac offers output that is at worst competitive with leading academic and professional modeling efforts for the Station scenario. Looking at the details of the other available modeling platforms, it is surprising that the results are so homogenous. Although SocEvac was not the most accurate platform when using a Level One analysis, it blossoms when Level Two when reasoning becomes an important part of modeling. SocEvac's simulation of Station nightclub fire based on the Level Two results is encouraging. While there is still significant room for future improvement, these results make sense, and they show that social behavior modeling has a promising future as one of several evacuation simulation modeling platforms.

Comparative analysis of results would show the importance of input reasoning (for example accurate buildings, population counts, hazards), reasoning (preventing end data from influencing simulation decisions), and output (the ability to do a Level Two analysis, not to mention reporting variability) are all key aspects when modeling real-world events. Although most model builders, including myself are often forced to include various caveats, we should at least set up explicit modeling frameworks so that one day these imperfections can be eliminated.

5.4 Level Three Analysis

While a Level Two analysis is usually sufficient to learn things about a building, but when trying to simulate behavior, we should go farther. The Level Two analysis above discusses a number of problems that arise from incorrect model inputs, incorrect output tasks, and overfitting of models. The Level Three analysis attempts to push evacuation modeling forward by following the individual fates of every avatar in the model, allowing for detailed studies that make it possible to identify and help solve the problems mentioned previously. The ability to do this is not new, for other modeling platforms have been able to log this information for years, but the reasoning for doing this level of analysis in one of the main contributions of this dissertation. An individual level analysis, tracking the outcome for every unique avatar, makes it possible to explore agent relationships, agent decision making, and model accuracy on a whole new level. In many ways, this is a textbook academic exercise that gives researchers the ability to explore problems in new ways, such as evaluating the validity of models in a manner not previously attempted.

5.4.1 Purpose of Level Three Data

The Level Three data makes it possible to use the results of the model to explore every possible element of the SocEvac model. By determining individual exit paths, group affiliations, evacuation choices, and other specific model parameters, it is possible to ask new sets of questions about evacuation models.

Level Three data tracks the model exits for every agent as individuals, meaning that the number of people that used an exit choice is not longer the end of model output. Using SocEvac and a good data set like that of the Station makes it possible to compare individual outcomes. For SocEvac, this individual tracking is important because it allows for analysis of which agents join or leave groups, which agents make individually suboptimal exit choices, and how many avatars in a model experience the same fates as their real life countparts. This level of detail makes location-based questions, outcome-based questions, and spatial regressions possible.

In SocEvac, where the main intention is to create a model that accounts for social behavior, Level Three data makes it possible to check not only whether social behavior functions are working, it also allows researchers to see how they are working.

5.4.2 Detailed Results

The Level Three results of SocEvac output is 464,535 lines of data, opposed to the 999 lines of data for Level One and Two results. Additionally, each line has 15 columns of data, opposed to one for Level One and seven for Level Two. This creates a data set that can be difficult to manage for higher numbers of model runs and more complex models than the Station.

The Level Three output returns the following variables: “Agent ID”, “Group Number”, “Energy Level”, “X Location”, “Y Location”, “X Goal”, “Y Goal”, “Following Behavior”, “Cause of Death” (if applicable), “Exit Used” (if applicable), “Orientation”, “Leader ID”, “Leadership status”, “Group status”, and “Supra Force”.

5.4.3 Comparison of Model to Actual Station Event

At the aggregate level, the results for Level Three are identical to the first two levels. However, the individual data makes it possible to compare the results of SocEvac to the actual Station event at a whole new level of detail. Instead of examining how many avatars use each exit, we can now compare the 452 (known) exit results to the exit results in each model run. This gives us the total number of individual outcomes that the model gets “right”. At this level of analysis, in order to be correct, the model must correctly account for a death if the real life counterpart perished, and if the real life counterpart survived the model must output the correct exit among six choices.

The histogram below in figure 5.9 shows the number of “correct” outcomes for each of the 999 runs of SocEvac. The model returns an average of 129 correct individual outcomes out of a possible 452 (28.5%). This means that on an average run, SocEvac returns accurate individual results for less than a third of the avatars. The IQR for the 999 runs is between 124 and 135, with a minimum of 103 and a maximum of 152 (both statistical outliers). Unlike the Level One and Level Two results, based on the 999 run sample, there is no possibility of this model returning even a result with all 452 outcomes correct.

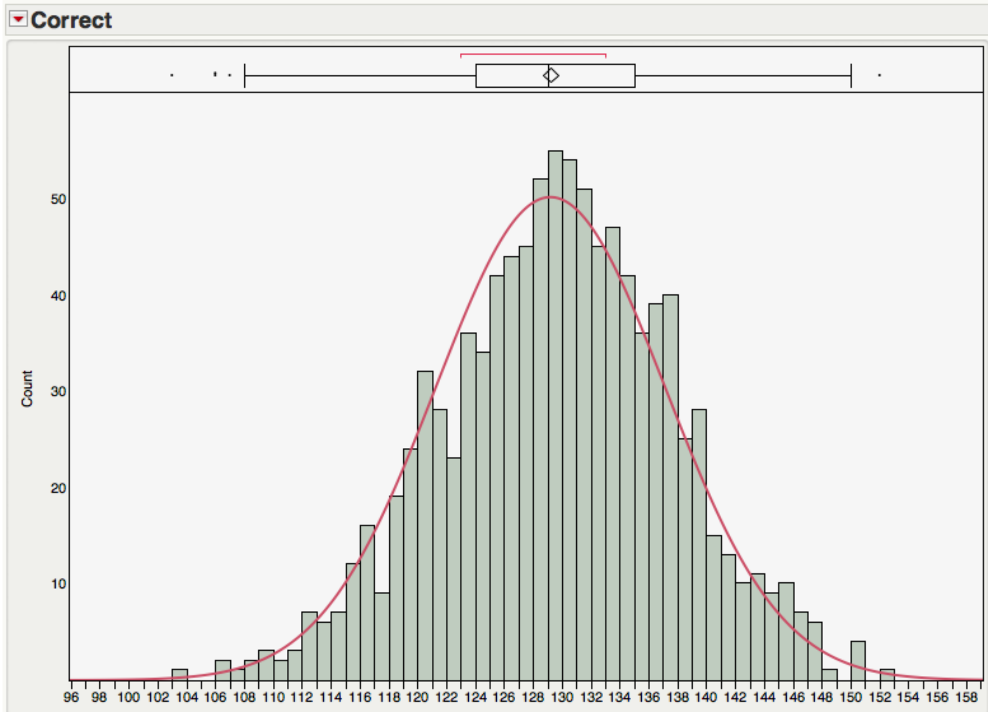

Figure 5.9: Distribution of Correct Outcomes for SocEvac

While it is unsurprising that all 452 known individual outcomes are not modeled accurately, there is still a great deal that can be learned from this data. Using Level Three data, we can figure out if there are meaningful differences in fates and demographics. Out of the 465 indivduals modeled, there were 75 that SocEvac never simulated correctly, consisting of 13 outcomes that were not verified, and 62 that were known but never correct. There were 6 outcomes that were correct in all 999 runs. This leaves 384 avatars that returned the correct result in SocEvac at least once but not every time. These are important numbers to track as high degrees of universal individual outcomes would indicate overfitting of a model. It is not surprising that there were more individual outcomes that were universally incorrect than those that were correct, since there are six wrong choices for 452 avatars and only one correct one.

There is a special interest in the 384 avatars that did not have a universal outcome in SocEvac (meaning they had at least two different outcomes in 999 runs). Essentially, avatars that have the same outcome 999 times are little different from their counterparts in closest-exit models with no variability. The most frequent non-universal outcomes were low levels of correctness, where the outcome was correct at least once, but less than about a dozen times. This indicates that the real-world results were certainly possible although unlikely according to SocEvac. On average, each avatar in SocEvac has the correct fate 278 out of 999 times, but the median is only 125, significantly lower than the mean because of the high number of zero and low correctness avatars.

Out of the 465 avatars in the scenario, there is a 100 avatar subset that should have perished in each model run based on the real event. Looking only at the 100 avatar subset representing those that perished in the actual event, over 999 runs of SocEvac this subset has a 41.6% correct fatality rate. While this percentage is far below 100%, it is well above (and significantly different from) the statistical likelihood of perishing in the fire (21.5%). Interestingly, it is disappointing that SocEvac returns such a low percentage given the fact that the program returns overall fatalities in the model were greater than in the actual event.

Of the 365 avatars that survived in the actual event, SocEvac has a 24.9% chance of correctly outputting the exit they actually used, and a 79.3% chance of correctly predicting survival. For the six possibility exit outcomes, SocEvac does a good job of predicting actual exit, but it is easy to see that the model has a long way to go before high levels of individual accuracy are achieved. Correct survival percentage is a different story. It is an accomplishment to return a platform that correctly replicates survival for almost four fifths of agents representing those who survived in the actual scenario. The fact that this is done without any exit assignment is a particular point of pride.

Merging these two subsets together, the model returns a 41.6% fatality rate for the avatars representing those that perished in the actual event, and a 20.7% fatality rate for the avatars representing those and survived in the actual event. This is a statistically significant difference that shows promise for this social evacuation model. The odds ratio for death in the model between avatars of those that perished and avatars of those that survived is 2.73.

When examining exact outcomes (correct exit or correct death), SocEvac returns a mean of 132 correct outcomes (max 452) over 999 model runs. However, when determining correct survival (correct death or survival), SocEvac vaults to an average of 331 correct fates (max 465) over the same 999 runs. This means that 71.2% of avatars have the same survival outcome as their real-life counterparts.

5.4.4 Comparison of Social Behavior Model to Closest Exit Calculation

As with the previous levels of analysis, there is a natural desire to compare the results from SocEvac to the results of a closest exit calculation for the example scenario. In the closest exit calculation of the Station scenario, 50 of those that perished would have perished in the model, and 67 survivors used the closest exit. In a run with no fatalities, 88 of the avatars had the closest exit equal their actual exit. As such, the closest exit model returns either 88 (only survivors) or 117 (survivors and deaths) out of 452 correct individual outcomes at a Level Three individual measurement.

When looking at the survival measurement, the closest exit model had better results, correctly predicting 31 of the deaths, and 341 survivors, for a total of 372 correct fates (80%).

Comparing these results to SocEvac we have an interesting disagreement. When comparing specific individual outcomes, SocEvac is the more accurate platform, but when comparing survival, the closest exit model is the slightly better predictor of the Station scenario. While SocEvac returns better detailed individual results, and better macro exit numbers (as discussed in the Level Two analysis), the closest exit platform is still a better fit when only looking at macro level survival percentage.

Much of this improvement in the survival numbers is due to the lower death count from the closest exit results. However, it does validate the closest-exit method for quick event reenactment on the macro scale, at least in the event of the Station. This again shows the importance of measurement. In a sense, the Level Three measurement of death or survival shows the closest exit calculation at its best, while the Level Three individual outcome result is highly lacking. This supports the idea that the reasoning for the closest exit calculations is inadequate, even if the results are sometimes good enough.

In essence, this data shows that being close to an exit makes you more likely to survive, but not necessarily likely to take the closest exit. Again, this means that the closest exit model might get the right results for the wrong reasons; the case against using these calculations is also strengthened by this detailed analysis of the results.

5.4.5 Guideline for Future Comparison to Other Models of the Station Event

While the comparison between SocEvac and the closest exit model suggests that the "best" model is not easy to conclude, it is unfortunate that at this point it is not possible to compare Level Three results for other models. At this time, replicating actual scenarios at a the individual level of detail is not a priority for many model builders, even though it can go a long way to helping explain how a model is working. For instance, if the closest exit model is returning better high level results but inferior detailed results, this should be the basis for a discussion about how to improve closest exit modeling, and should not serve as validation that closest exit models are superior. Ideally, other model builders will make Level Three data available in the future so that model users can determine how models are functioning in addition to the results that they output.

The Level Two analysis in the previous subsection provides a justification for examining output data in more detail is appropriate, especially when modeling actual scenarios where detailed results are known.

5.4.6 Level Three Conclusion

This Level Three data suggests that the SocEvac platform has the ability to return highly accurate results in the Station scenario, and these results may inspire model builders to consider adding social behavior elements and detailed inputs into future evacuation models, not to mention making it possible to track individual level results.

When looking at individual level data and specific individual outcomes, the social behavior model shines, returning results that show the promise in the platform, but in a comparison to closest-exit models, the aggregate death and survival ratios show that there is a role for closest-exit modeling in simulation, even if it is only to be used as a baseline.

These results are fascinating in that how you measure determines what is the best model. As model builders, we need to decide if we are more interested in aggregate or specific outcomes. Although the closest exit platform offers an improved aggregate survival and death number when examining individual death and survival, it does so in a manner that returns inferior results at all three levels of measurement.

Ideally, the “best” modeling platform would return results are superior across the board, with closely fitting results for all three levels in addition to a good fit for model versus actual outcomes. While SocEvac is not yet accurate enough to outperform competitors at every level of analysis, these results show that social evacuation modeling should at least become part of evacuation modeling discussions, since it can be shown to outperform other modeling platforms including the closest exit model when examining individual fates (and isolated exits).

5.5 Beyond Level Three – Detailed Intra-Model Data

While the Level Three analysis is already farther than most researchers will be willing to go when testing models, SocEvac outputs levels of data that allow for analysis far beyond the Level Three results. It is possible for SocEvac to track all of the Level Three data for every avatar up to 20 times a second. This would result in 2,418,000 rows of data for one simulation (more than the 464,535 rows from the entire 999 run Level Three analysis). At this level of detail over 999 runs, SocEvac could return a 2,415,582,000 row dataset with over 20 variables in each row, creating an output file that would be almost unmanageable with consumer systems. Because gathering data about slow moving pedestrians is not required 20 times a second. To make analysis easier, data was collected once every 15 seconds hoping to create more manageable files. The result was a large 8,361,630 row data file that allows for the examination of things like group identification, group adoption, group abandonment, leadership changes, and supra force. Because the end times of the simulations can actually be very different, this analysis will look at the first 11 measurements, taken at 0, 15, 30, 45, 60, 75, 90, 105, 120, 135, and 150 seconds in simulation time. After this point, the numbers become distorted because some simulations are finished. In order to create data file that can be used on common analysis platforms a randomly selected 100 run model sample was used for this analysis.

5.5.1 Group Interaction and Leadership

While group membership is one of the most important differences between SocEvac and other modeling platforms of the Station, examining the results from the end of the model only do not allow for a clear picture of the role that groups play in the evacuation. In order to figure out when groups are forming, we need to be able to see results over time. The best way to do this is to visually watch the model, because it is more intuitive to watch groups actually come together, but it is also possible to use the SocEvac output files to see what is going on in the model.

Looking at the results in 15-second intervals over the model run, there are 179 leaders at the inception of SocEvac, including agents that are acting individually. This number is a constant because at inception no avatar has left the building environment in the model. As leader agents exit, a new leader would be selected from the remaining group members, or if agents were acting alone as leaders, there would be no leader taking over. Because leading agents acting individually took some of the shortest paths to exits because they experience no “inefficient” pathfinding group behavior, agents acting as individuals are more likely that group members to exit the scenario quickly and to survive.

Before analyzing these results, a review of the social groups in the Station scenario is warranted. In the DRC analysis of the Station, there were a total of 48 individuals that were in the building without social bonds, and 417 agents that were in social groups with two or more members. There were 68 groups of two, 25 groups of three, 19 groups of four, nine groups of five, nine groups with six to eleven members, and the largest social group, the employees, with 17 members. Because of the way the SocEvac functions work, after accounting for other factors, the larger a group is, the less likely it is to disband. This is because small groups are more likely to have one or more agents exit or perish in the model, and agents of small groups are more likely to be farther away from their closest group member, assuming even distributions of agents.

After accounting for the leaders, there are 286 potential following agents that can select to follow their group leader or evacuate as an individual agent. Of these 286 potential followers, 175 agents are in 2-4 avatar groups, while the remaining 111 are in larger groups. It is important to remember that a group is abandoned when there is only one member left, but that individual abandonment of larger groups does not dissolve the group itself unless all agents switch to exiting individually.

At the inception of the model, an average of 144 agents are following another leader. This means that 50.3% of agents that have the potential to follow another are choosing to do so at the start of an average model run. While this seems to be a low following percentage, it makes sense when reconsidering the programming of SocEvac. To begin, there are many avatars that have one or more of the conditions that make it likely to abandon a group: close proximity to an exit, no close proximity to other group members, and loose social bonds. Based on the high distribution of avatars that are far away from their group members and close to an exit (such as a male dating partner in the bar area and a female dating partner on the opposite end of the building by the stage), it is reasonable to assume that individual exit proximity would overrule the desire to evacuate with a group member in all but the tightest social bonds.

By 15 seconds into the simulations, an average of 85 agents are following a group leader. This means that about 59 avatars abandon their group, have a group size of one, or exit the simulation in the first 15 seconds. At 30 seconds, the average number of followers is 54. At 45 seconds, the average number of followers is 42. By one minute into the simulation, the average number of followers is reduced to 33. In the next four 15-second intervals, the average number of followers are 23, 16, 11, and 6. By 135 and 150 seconds, the average number of followers are 2 and 0. This makes sense within the quickly lethal Station scenario.

Based on these results, SocEvac is functioning exactly as intended as far as tracking social bonds in an extreme hazard environment. While a number of agents are willing to abandon their groups at the onset, many agents choose to remain in their social groups even after the building environment has deteriorated significantly and survival is called into question. These numbers coincide fairly with the distribution of group types in the Station. Obviously, the more agents that exit a building or perish, the smaller the number of active groups will be. By 150 seconds, active agents would have significant damage from smoke, and it makes sense that any active agents would abandon their group at that point in order to attempt an individually optimal exit.

5.5.2 Fatalities Over Time

In the Station scenario, there are two causes of death, fire and smoke inhalation. As discussed previously, fire and smoke damage occur differently in the program. If an agent is surrounded by fire, death in the simulation is almost immediate. Smoke damage occurs over a longer period of time, and depends on several factors, including the demographics of the agent and the location of agent (there were specific areas in the Station, mostly near doors and windows, where studies showed that survival was possible for longer periods of time) (Gill et al. 2010). Analyzing average model fatalities over time shows the dual hazard function in the Station scenario is working as designed.

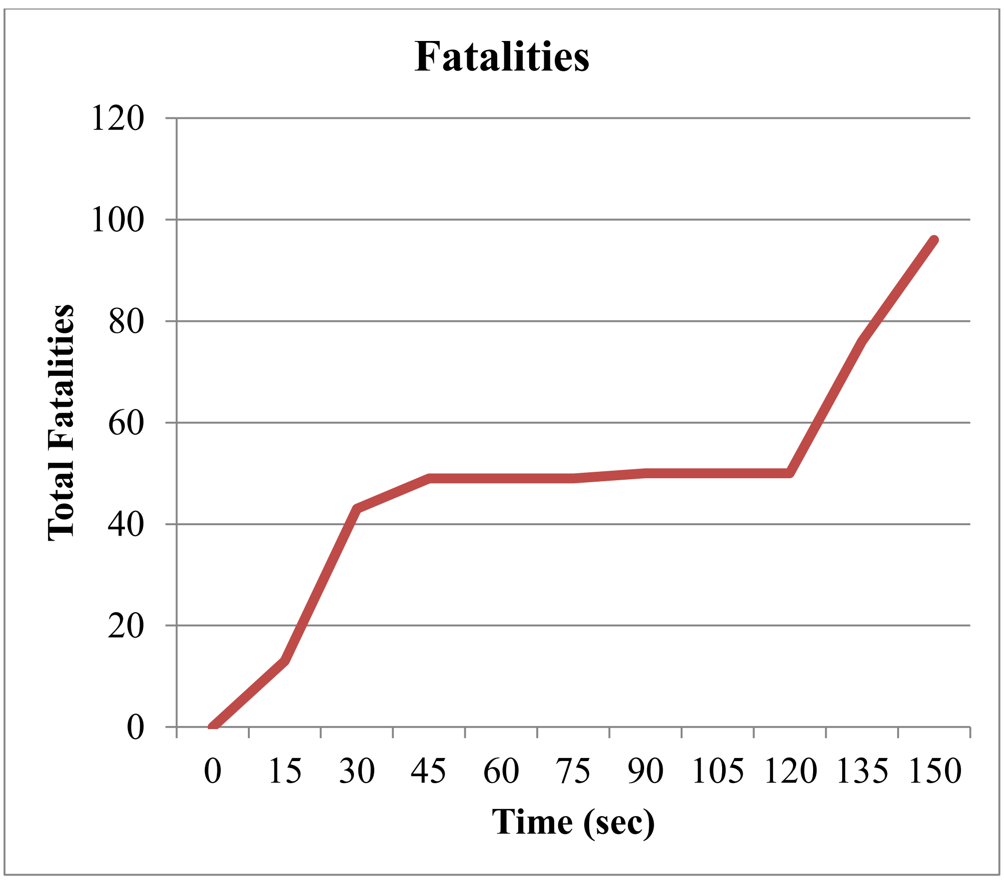

Figure 5.10: Model Fatalities over Time

Figure 5.10 shows the average fatalities every 15 seconds from model initialization to 150 seconds of simulation time. Within the first 45 seconds, a mean of 49 avatars succumb primarily due to the fire. This death total remains relatively constant until around 120 seconds when avatars begin to perish from smoke damage in large numbers.

Many of the other models of the Station assume relatively linear exit and death numbers over time, but SocEvac in its current form has a significant difference between the mean time of fire deaths and the mean time of smoke deaths. SocEvac damage was designed based on the findings of the Rhode Island medical examiner after an analysis of the Station tragedy. This included the separate fire and smoke damage profiles (with smoke damage being a composite of multiple different gasses), and the concept that all areas in the building are not equal. Although the smoke is potentially fatal long before 120 seconds in certain areas of the Station, most of the avatars have either left these areas or succumbed to fire damage within the first 45 seconds. At 60 seconds and beyond, most of the remaining avatars are largely concentrated in the atrium area leading to the main exit of the bar area. These areas are in close proximity to exits and windows, making these areas more insulated from smoke than the remainder the nightclub.

The average number of fatalities at 150 seconds is much lower than the average number of fatalities at the end of the 999 measured simulations. This is because results were cut off for the 100 model sample after some simulations stopped reporting, which would bias the time-series analysis. There appears to be a correlation between simulation time and number of agent fatalities.

The next step was to determine if there were differences in the incidence of death by demographics. In a situation like the Station where seconds matter between life and death, there are significant differences in time spent in the environment based on demographics based on a number of factors specific to the Station, such as the distributed location of group members, and the different balances of genders in the stage and bar areas. Death incidence was measured over time by familiarity with the environment, gender, group size, and group type.

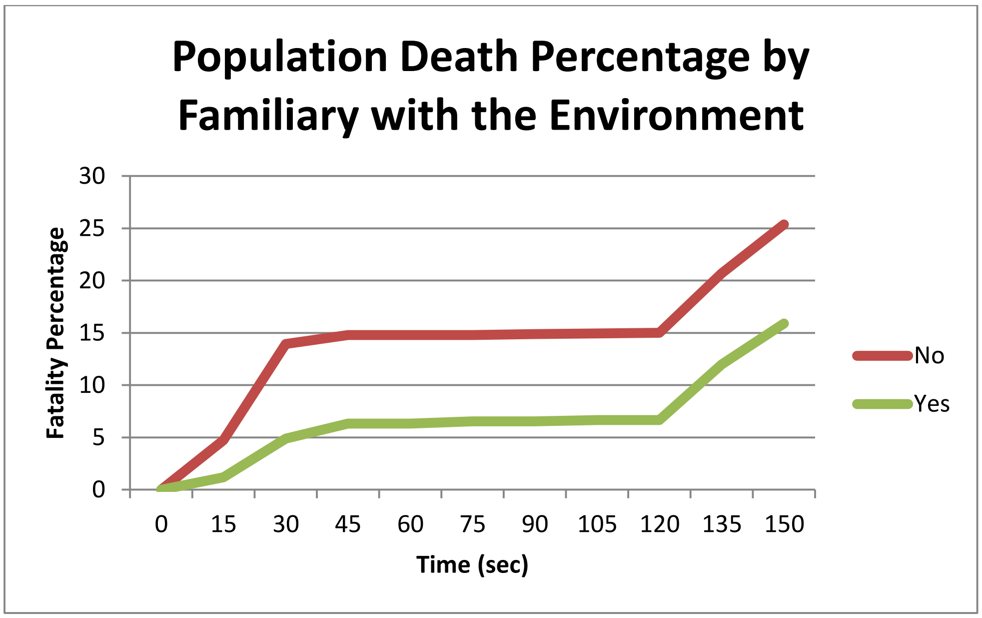

Figure 5.11: Death Rate by Familiarity

Figure 5.11 shows fatality percentages for subpopulations divided by familiarity with the environment. It is easy to that that those avatars that were familiar with the environment experienced lower fatality levels throughout simulation runs. This makes sense in both the context of the scenario and the model, because those avatars familiar with the environment were more likely be aware of multiple exit paths.

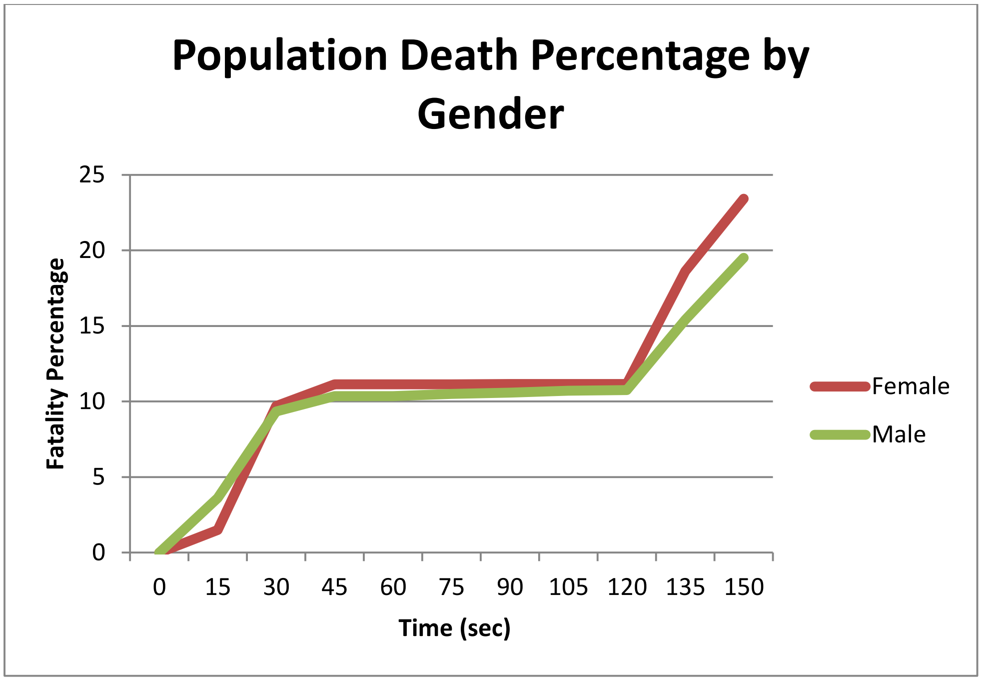

Figure 5.12: Death Rate by Gender

The next subpopulations measured were the male and female gender. There is a small difference in fatality percentages between genders mostly at the end of the simulation. When comparing these numbers to the closest exit model, it appears that in this single scenario, these differences may be more of a function of starting position (many females were in the stage area and could not quickly evacuate) than gender itself. It is not possible to make any conclusions about gender in this case.

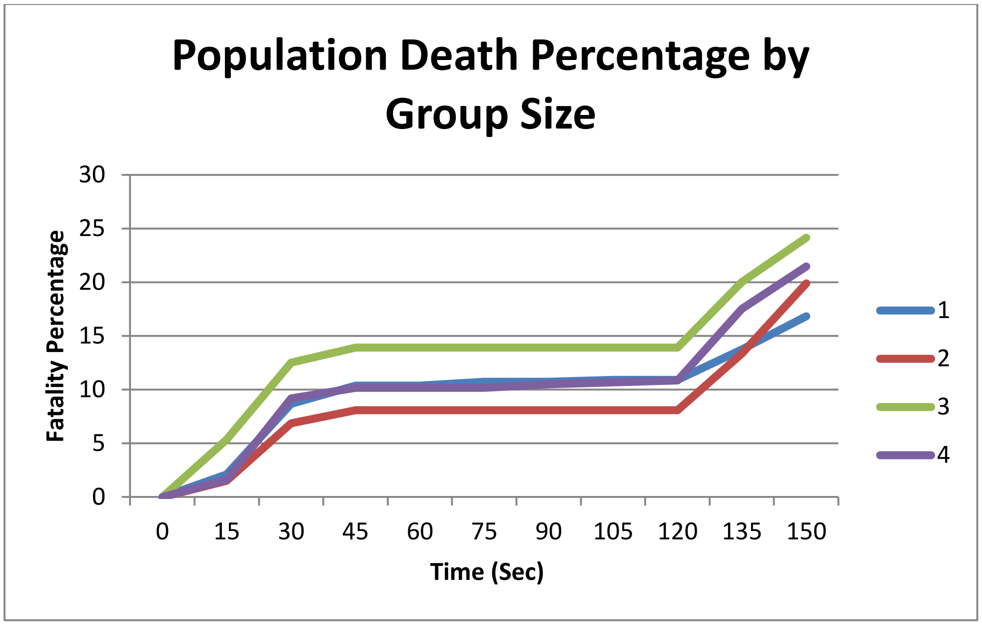

Figure 5.13: Death Rate by Group Size

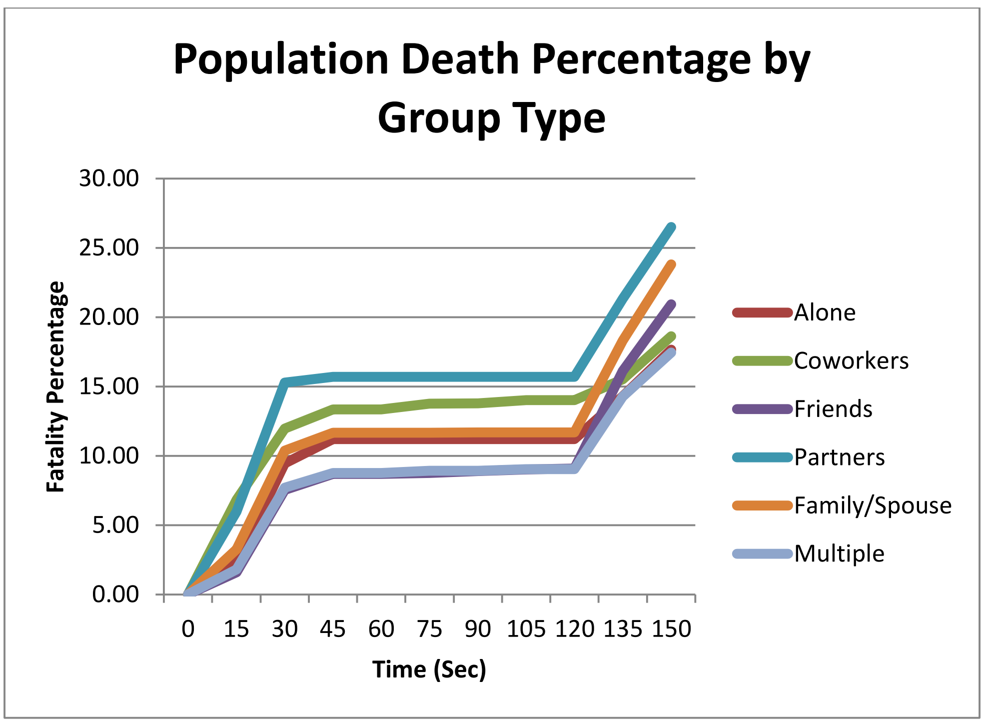

Figure 5.14: Death Rate by Group Type

The two figures above show fatality percentages for each subpopulation isolated by group size and group type. While these both confirm that individuals generally have lower fatality percentages than group members, this phenomenon is only true towards the end of simulation runs. While both of these factors generally make fatality in the model more likely, they are far from absolute. This is likely because of small sample sizes for subpopulations in the scenario, and because positioning in the Station does play a very important role in potential fatality percentage.

While avatars certainly do not always take the closest exit, being near an exit is a trait that helps reduce potential fatality. In the future, it would be desirable to model this scenario and other scenarios with spatial regression to determine the relative importance of positioning versus demographics and group associations.

5.5.3 Supra Crowd Force

One of the major contributions of the models developed at the DRC to model to model the Station nightclub scenario is the inclusion of "supra force", or the concept of environmental forces that override an agents' ability to move as they desire. In an environment like the Station, where crowd density was a major impediment to evacuation, it makes sense that many agents would become surrounded by a barrier of other agents, especially when agents are no longer all taking individually optimal paths because they are attempting to evacuate with other group members.

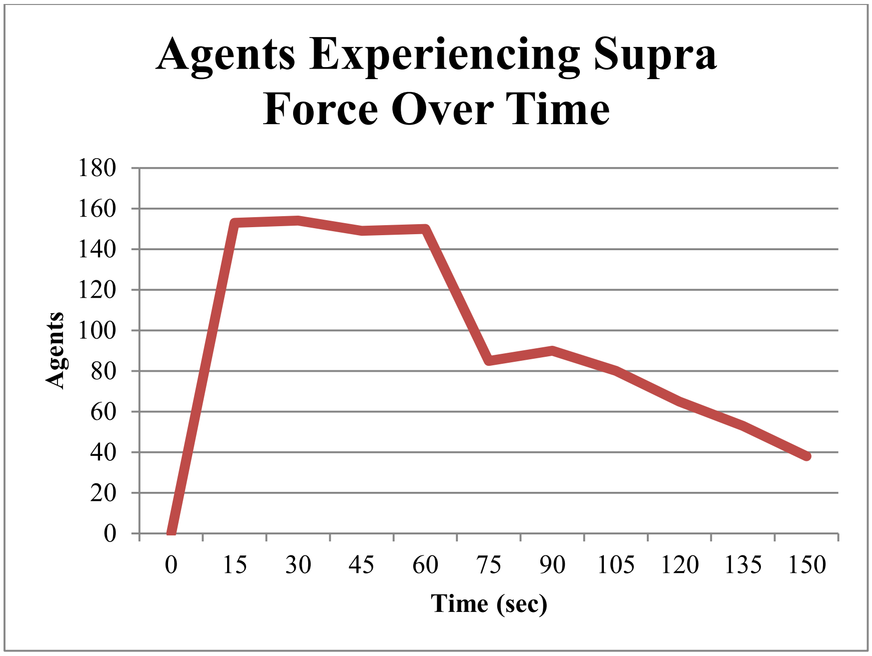

Figure 5.15: Agents Experiencing Supra Force Over Time

The graph above shows the importance of supra force in the Station scenario, and goes a long way to explaining why there are so few following agents even 15 seconds in the simulation. From 15 to 60 seconds, the average number of agents experiencing supra force is between 149 and 154, close a third of all agents at the initialization of the simulation. This is a significant percentage of the simulation population, and it highlights both the importance of including this metric in the model and the difficulties to evacuation created by social behavior modeling.

At 75 seconds and beyond the number of agents experiencing supra force drops sharply from the levels seen in the first minute, however the rate of reduction slows later in the simulation. This is because there are still many agents that are part of strong social groups that are refusing to take individually optimal paths, causing blockages for others.

The next logical analysis was to determine supra force incidence by demographic. Based on the way the program was coded, it would be expected that certain demographics more likely to cause contraflows would be more likely to experience supra force. In order to test this, the DRC measured supra force incidence by familiarity with the environment, gender, group size, and group type.

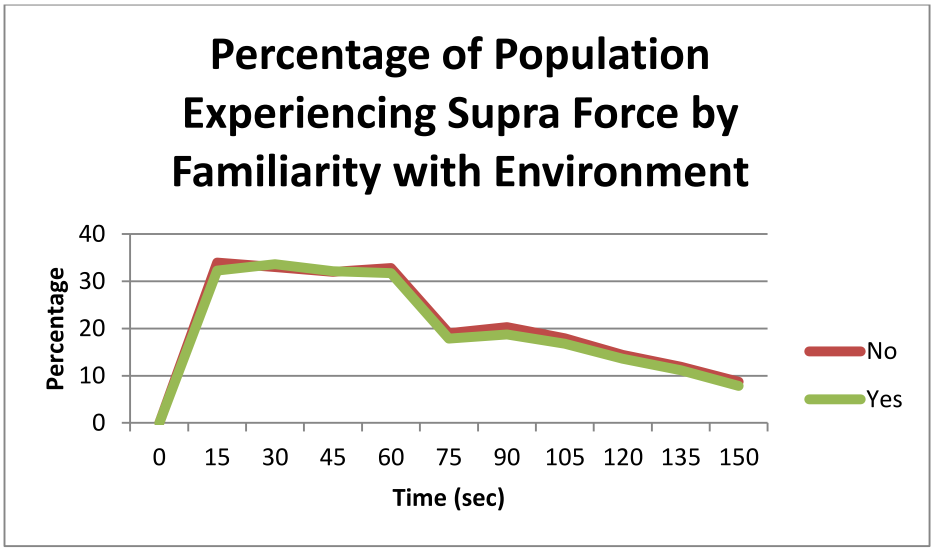

Figure 5.16: Supra Force Rate by Familiarity

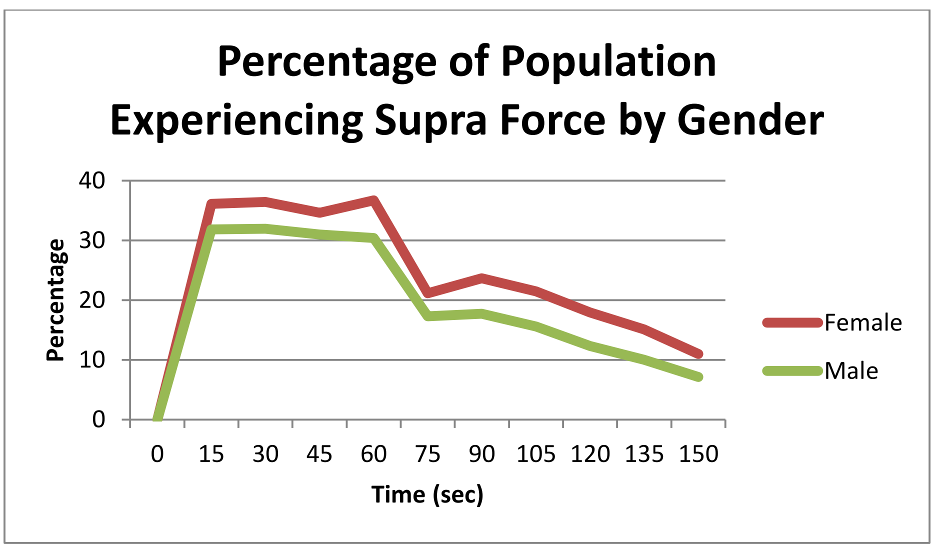

Figure 5.17: Supra Force Rate by Gender

Figures 5.16 and 5.17 highlight supra force experience for the familiarity and gender subpopulations. Surprisingly, familiarity with the environment did not significantly reduce the likelihood of experiencing supra force. Looking at gender, females experience higher levels of supra force in the Station scenario. This appears to occur because a higher percentage of females are followers in intimate groups, and these avatars get stuck in large throngs of evacuees trying to find their group leaders.

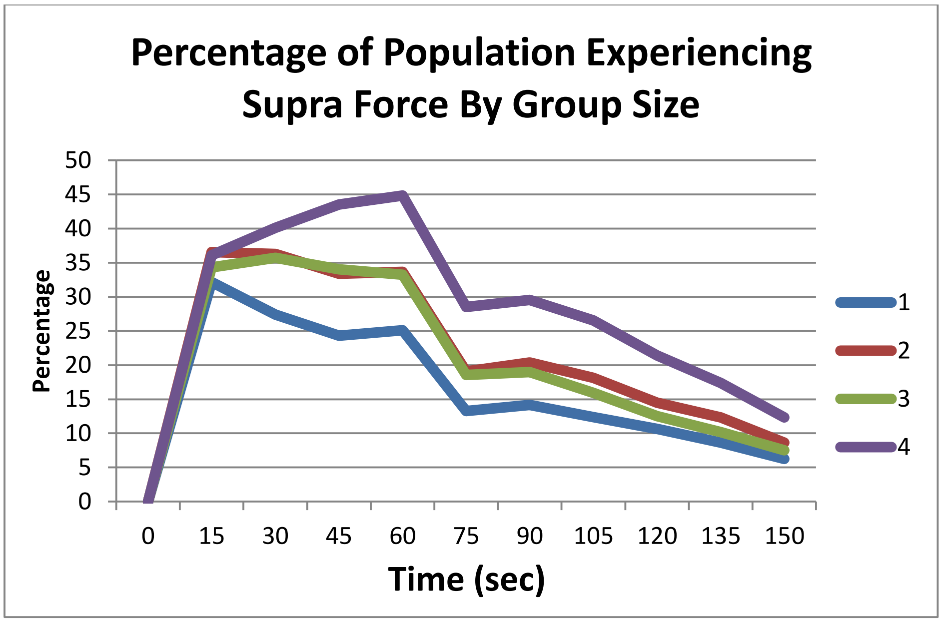

Figure 5.18: Supra Force Rate by Group Size

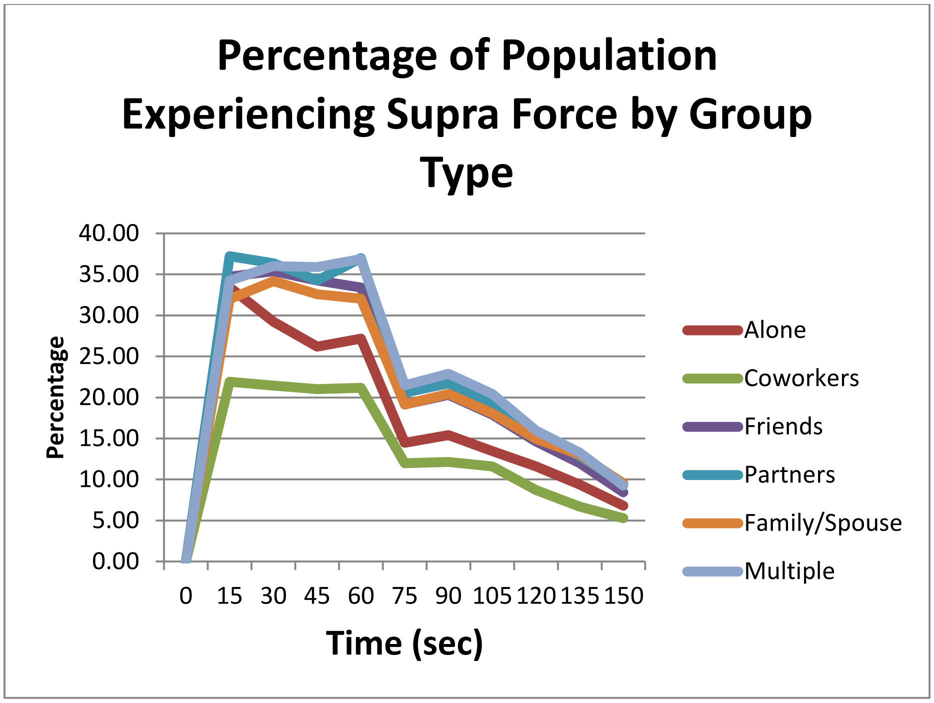

Figure 5.19: Supra Force Rate by Group Type

Figures 5.18 and 5.19 isolate subpopulations by group size and group type. , As expected, being a member of a large or intimate group generally makes avatars more likely to experience supra force.

As these numbers suggest, group abandonment in SocEvac is likely less a function of agent choice, and more a function of crowd density overriding the decisions of individual agents. Obviously, in the Station scenario, supra force is a significant part of the barrier that agents face in evacuation, which conforms to the DRC studies about the Station nightclub fire. While it is beyond the scope of this dissertation, supra force should be included in all models of building evacuation in addition to group behavior, as these two elements are interlinked. Contraflows appear to be vitally important to accurate simulation of building evacuations.

5.6 Alternative SocEvac Output

The SocEvac experiment in this dissertation focused on reasoning, not results, but testing done to the model to ensure correct reasoning frequently resulted in changes to the model that provided more accurate (and in some cases highly overfit) results. While these results were not included in the comparison to other models, it is worth noting that the platform can sometimes obtain better results by loosening the rules about reasoning (or more nefariously, targeting a survival percentage).

In addition to the normal or “medium” damage profile based off the best data about the Station, the DRC ran 1,000 run simulations the Station scenario with both a higher level of damage and a lower level.

Looking at these results from a Level One vantage, the fatality levels decreased for the both the lower damage and higher damage models. While that was initially surprising, looking at the results in more detail revealed some interesting correlations. The lower damage model gave avatars more time to evacuate without perishing due to smoke damage, but the higher damage model reduced density more quickly, allowing more efficient evacuation (dead avatars do not block egress of others that same way active avatars do). Both alternate versions of the model produced more accurate fatality levels, averaging 94 and 110 deaths respectively. Deaths from the light damage profile ranged from 40 to 245 and deaths from the high damage model ranged from 51 to 278. The Level Two results were varied, but the Level Three results were almost identical in averages of correct fates to the damage profile chosen for the main analysis.

For further analysis in this field when more processing time is available, it is my hope to use this same scenario to analyze the effect of small changes in all variables in a standardized way. Non-standardized results tracking during model development were not included because there is no way to verify if output differences are the result of variable changes or other coding changes as the model developed over time. Based on the analysis of these three 1,000 run batches, this exercise would be a valuable addition to a mature modeling platform.

5.7 Meaning of Model Results

While adding the elements that make SocEvac unique among evacuation models led to clear differences in results, these results were not always "better". SocEvac returns plausible results upon examination of analysis at three levels discussed in this dissertation, in addition to providing output data that withstands scrutiny when looking at evacuations over time.

While it was a valid criticism to assume that these improvements in results may not have been worth the effort, now that this platform, and a few competing models, are available, the argument will hopefully shift from rather it is worth it to program these attributes to whether they should be included in more micro-level simulations.

5.7.1 Which Level of Results are the Most Accurate?

Based on the initial Level One analysis, SocEvac does not look like any kind of revolution in building evacuation modeling. After lots of extra effort, SocEvac returns results for the Station scenario that are no less bad than competing platforms. This is not a surprise since it is easy to tune models to get overall death count very accurate, and modifying a program to fit results was never the intention of SocEvac. Looking at results at this level of detail suggest that social behavior does not make anything worse, but it does not necessarily provide any improvements to a model of the Station fire.

Level Two results are where SocEvac begins to show significant differences from peer models of the Station scenario. These results are a combination of both extra effort put into accurate modeling of inputs and the benefit afforded by the social behavior model. While all models have a difficult time correctly mapping Level Two results, it is fair to make the case that SocEvac is one of the best choices based on the comparison to model results to the actual identified use of exits. The Level Two results prove the worth of social behavior modeling, and suggest that the topic merits continued exploration.

Level Three results also indicate that social behavior modeling might be worth the effort. A clear majority of Station occupants did not use what the DRC identified as their closest exit. While SocEvac also falls far short of perfect mapping of model to reality, there are improvements afforded by the model when predicting specific fates, although these benefits are not as pronounced when only measuring survival. While the Level Three results make another strong case for the inclusion of social behavior, the reality is that this level of detail will be unavailable in most circumstances.

Regardless of the level of analysis used, SocEvac returns results that make sense, and the platform has value being added to a larger set of building evacuation models.

5.8 Results Conclusion

Level Three results should be the baseline for modeling when Level Three data is available. However, Level Three data is only available in situations where significant resources are devoted to analyzing behavior after an emergency. In most cases, Level Three data will never be collected, and it is probably overkill to consider this level of detail in most models, including all predictive models. This is even truer for the data that is more detailed, such as the time series data for individual behavior. While this highly detailed data allows for deep exploration of what is happening in a model, it may be more important for model development than analysis of model results. As mentioned in previous sections, it is possible to generate so much data from these output files that it becomes difficult to analyze without commercial-level computing systems, resources that are not available to many potential model users.

In most cases, Level Two analysis is the best combination of potential accuracy and practically. Using the results from this dissertation as an example, the Level Two data allow a user to see a clear difference between a social behavior model and a closest exit model without requiring details to be known about all model participants. By analyzing exit choices, it is still possible to learn a great deal about how models function without requiring the expensive and time-consuming data gathering to track individual fates. The Level Two analysis in this dissertation shows the differences (and similarities) between various modeling platforms, suggesting that benchmarking may cause competing modeling platforms to return very similar results.

Based on the Level Two results, model builders should consider SocEvac and the concepts it attempts to address in future micro-level modeling efforts. Adding social behavior is more an issue of convincing model builders and users of value than it is of capability, and the results highlighted in this chapter make a strong case for addition of social behavior concepts to models.

Chapter 6

FUTURE IMPLICATIONS AND LIMITATIONS

While we have a surplus of models that return reasonably accurate macro results, models that are completely accurate at the micro-level are likely a long way off. SocEvac is an attempt to help with the reasoning side of the equation, even if the results from the model are not always superior. In order to make models better, we will continually need to begin by pushing boundaries with specialized models that by definition are not appropriate as generalized platforms. As these specialized models are refined and advanced, they can be incorporated into more general platforms. Having been through this process with the first-generation social evacuation model and then SocEvac, there are benefits to seeing projects through from prototype to general-development phases, and I hope to continue involvement in agent-based modeling of building evacuations beyond my graduate career.

There will always be new advances in computing power, coding methods, and data collection, but we cannot benefit from these fully until we understand what exactly we are trying to recreate. Ideally, building evacuation models would be used to save lives, and in order to see this reality, we need to create platforms return correct survival percentages through accurate reasoning, making it possible to test the “what if” questions that will continue to arise. While the solution offered in this dissertation is a step in the right direction, and it goes a long way to addressing the challenge presented by previous DRC researchers to incorporate group behavior, there is still much progress required before we can confidently predict accurate micro-level behavior for a multitude of environments (Torres 2010). In a sense, this dissertation offers its own challenge; after showing that it is indeed possible to use social behavior and supra force equations to return a micro-level model that offers some advancements on closest-exit or fastest-clearance models, the modeling community needs to figure out how to incorporate this type in evacuations designed to recreate situations with social relationships. As mentioned in chapter 5, this is no longer a discussion about the capability, but rather one about understanding that these relationships are significant enough to effect results.

6.1 How Detailed is too Detailed?

This dissertation only scratches the surface of what is possible with models accounting for both social behavior and supra force in addition to individual behavior, but it is important to include all three of these elements in simulations, because the rules for allowing social behavior in evacuation models create the conditions for supra force, and based on this first test scenario, all three levels of possible behavior are used by a significant number of agents.

There is a discussion to be had about the detail of both model input and output. Beginning with model inputs, while these results are promising, there is much more that could be done to ever more accurately model building evacuations. The most important decision is to decide where to stop in the name of speed, efficiency, and diminished returns. Based on the experience doing the work that led to this dissertation, including social behavior and supra force in some form are complications that are absolutely worth adding to micro-level simulation models. It was not possible to recreate the results of the Station scenario without using social behavior elements. While this single case is not a justification to fundamentally change evacuation models, there is a substantial body of literature that supports the idea of social evacuation.

While there is some arcane data analysis in chapter 5, the results included in this dissertation only scratch the surface of the types of analysis possible with Level 3 data and detailed data about model runs over time. My ability to analyze these results even became limited by current storage and computer processing power. For more complex simulations, such as skyscrapers or stadiums, we will need to make significant advances in processing power both to run simulations in a timely manner, and to evaluate results. While we can obviously do more to analyze model results, it is possible to examine models in too much detail, glossing over larger and likely more important trends.

Level 1 results should no longer be used to determine the validity of building evacuation models. While the aggregate fatality levels of models should be reported, this single number should not serve as a standard of whether a model works well or not. A more reasonable approach is to examine both Level 2 and Level 3 results, and a combination of these two types of output should serve as the standard for which models are judged. Further out, after more detailed model analysis becomes mainstream, the results of intra-model runs should be included to determine that model reasoning is correct. While this process is many years away in the mainstream of model development, it should be the goal that model builders work towards, because model validity is really not possible if researchers are unable to determine what is happening during a model in addition to the end results.

While SocEvac is one of the first efforts to include social behavior in evacuation, there are other platforms that have surfaced that also address the concept. Thankfully, the landscape of model builders interested in adding social behavior to models looks much different in 2013 than it did in 2009. Interest in social behavior modeling continues to grow as improved results come in from current models and customers demand more accurate simulations.

6.2 Limitations

While incorporating social behavior is a major step forward in these models, there are a number of things that current social behavior models do not and will not address. While SocEvac and some peer models now account for much more group behavior than previous efforts, there are still countless scenarios that these current platforms cannot hope to address. Obvious limitations are the lack of emergent behaviors and groups, and the inability to create positive models that account for any foreseeable hazard. Usually, we first have to experience a specific building hazard before we are able to model for it, so while models may get better at assisting emergency managers with keeping gatherings safe, they fall far short of predicting which hazards will occur and how they will influence populations.

Specific to SocEvac, I hope to continue refining the code beyond my dissertation to include emergent groups, and to include much more flexibility for crowd interactions in addition to building ever more complex avatars. I hope to specifically add abilities for agents to communicate with each other and share advanced information, an ability that is now available in some building evacuation modeling platforms. This will allow for another level of realism in simulations, and the industry can move towards models that allow for both advanced information sharing and social group behavior.

While SocEvac offers some valuable insight into potential improvements for a widely modeled scenario such as the Station, there are number of modeling specifics that will be preliminary until more data sets as detailed as that of the Station become available. For instance, leadership equations, group abandonment percentages, group behavior, and the advantage of prior knowledge of an environment are all coded in SocEvac based primarily on the single occurrence of the Station evacuation. This is easily the biggest limitation of the model, although it was still worth it to include them in a model and test their importance.

6.3 The Future of Social Behavior Evacuation Modeling

While there is a great degree of diversity of focus among these model builders, it is promising to see a convergence towards the inclusion of more advanced individual and social behavior of simulated evacuees. As mentioned previously, when we began this project in 2009, there was not very much attention given to adding group interaction into models. Today, even though the obstacles to adoption in many models are still significant, there is widespread interest in including these features, and there are a number of models that now include some form of interaction. The effort described in this dissertation is one of these platforms.

To facilitate this effort, I have worked with other model builders over the course of this dissertation, sharing some of the code from SocEvac, in order to encourage the inclusion of social behavior models. This will lead to improved models and to more situations to test the validity of the conclusions from my work with SocEvac and the Station. Hopefully, in another four years the SocEvac engine will be one of several pseudo standards for adding social group behavior to evacuation models.

6.3.1 Future Data Gathering and Validity Testing

One of the largest barriers to making a model with the detail of SocEvac better is access to data. While the DRC attempts to gather as much data as possible about disasters including building evacuations, for a number of reasons the Station fire was by far the most complete evacuation dataset available. In order to get better output, we need more quality inputs.

In 2011, Apple Corporation created a mobile phone upgrade for their iOS platform that accidentally backed up private mobile device data on computers synced to users' iPhones (Apple 2011). These files contained troves of data about device users, and provided a window in future input data possibilities for programs like SocEvac. This mistake by Apple and AT&T allowed the civilian research community to see exactly what kind of data are kept by mobile carriers, and it suggests new possibilities for studying disaster preparation, response, and recovery. Real-time data is promising on its own, giving disaster responders and researchers quick potential snapshots of communities, but the real power is in the data sets that measure multiple points over time. Using these data for a community of users, researchers would be able to see what percentage of a population was currently in locations in and out of the ordinary. This would allow for analysis of population displacement, traffic patterns, evacuees, and people in danger. From an academic standpoint, these data also have the potential to change the nature of the disaster survey, allowing for targeted surveying and response verification in addition to data for modeling.

Following instructions published online (Warden 2011), I accessed my personal iPhone location tracking files, and after a few hours of searching other files backed up from the phone, had a wealth of data about my whereabouts. There were about 46,000 data points from a ten-month period. About 9,500 observations were from cellular radio locations, and about 35,500 were from Wifi locations. These data are in addition to stored GPS location points.

These values are stored in a way that makes seconds distinguishable, but unfortunately they are truncated to fewer time points per day. WifiLocation data is stored in decimal degrees to two decimal places, which distinguishes points to about 1.11km or less at common United States latitudes. CellLocation data is stored in decimal degrees to eight decimal places, which distinguishes points to 1.11mm or less. Looking at these data and at current technology, it is not possible to obtain accuracy levels beyond about four to five decimal degrees. The data are not possibly accurate to the millimeter, but the storage space is allocated in the files indicating hopes for technological improvements in the near future.

When data are examined by time and by location they become much more powerful and predictive. Looking at time and location together it was relatively easy to determine which times and locations had clear and consistent patterns. In the early morning hours the phone was almost constantly in a residence, during weekdays and work hours the phone was usually in an office. Overall it became possible to confidently predict where the phone would be based on day and time about two thirds of the time. Weekday evenings and weekends were difficult to predict based on a preliminary analysis, but nighttime hours and the work-week frequently returned the same locations. At this time it is possible to determine if my phone is “off pattern” more than half of the time based on past data. By using data from a large number of users with different predictive times, it would be possible to determine whether an area or time had a larger than expected number of users off pattern possibly indicating some type of anomaly. This could be done without contacting the users, without looking at their location, and without singling out any one user.

At this time, these data points collected by Apple and AT&T on the device itself are not detailed enough to track a user by the minute, but the collection methods are clearly set up in ways to suggest this ability is desired in the future, and AT&T suggests they can do it (AT&T 2012). While these data are invasive to track for one user, they can be useful in the aggregate without similar privacy violations. In the field of disaster science, data like these from a multitude of users could be studied to determine things like who responds to evacuation warnings, where people congregate in disasters or potential disasters, and many other purposes including using these data after disasters as inputs for far more detailed social models.

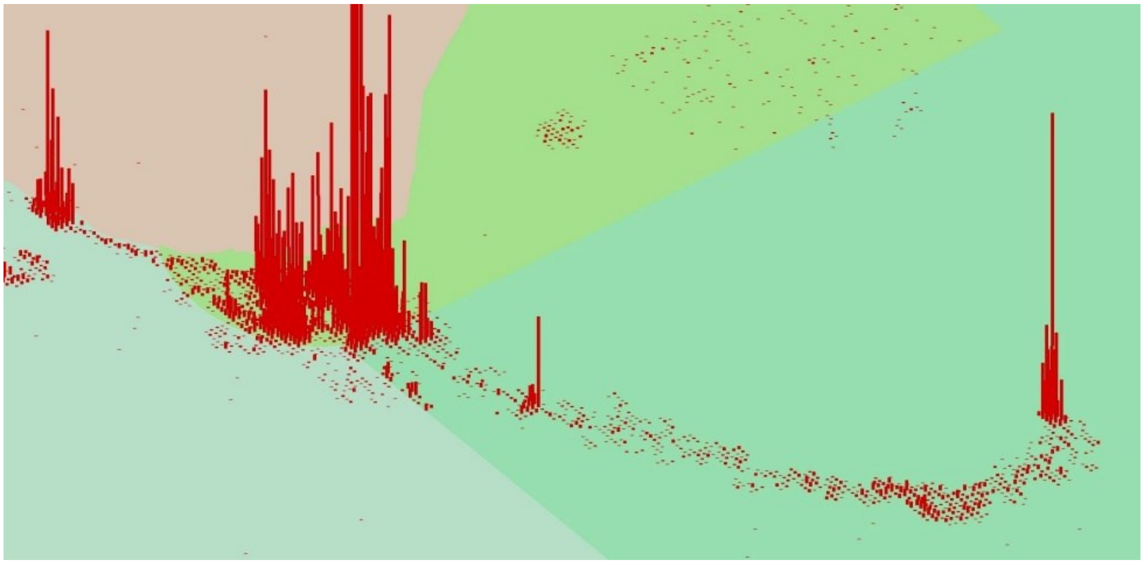

Figure 6.1: Aggregate Location Data for Eric Best's Mobile Phone

Figure 6.1 shows location data along the I-95 corridor for my phone under the sample observation period. The map shows Pennsylvania, Maryland, Delaware, and New Jersey with the Philadelphia airport in the upper left and Annapolis, Maryland in the lower right. X and Y points are location, and Z values (height) are the frequency of tracking for each location over a year.

With data from multiple users, relationships would be easy to define, and privacy could be invaded to an even greater extent while making identity confirmations faster and more robust as evidenced by past research (Calabrese et al. 2011b; Zhang et al. 2011). Merged data sets are becoming so complete that it is even possible to examine subject emotion by categorizing message content in addition to behavior or location (Thelwall et al. 2010). As more devices come online and more websites track user data, these data sets will only get more comprehensive and the results more accurate and powerful.

This level of detail, when made available for large groups of people, will allow researchers to validate models in ways that have previously been unimaginable. Connected device data will be the future of disaster after-action reporting, and the same data will lay the foundation for the next generation of behavioral modeling. By determining who contacts each other on mobile devices coupled with saved location data, it will be possible to obtain accurate data inputs for models of cities or even countries with enough processing power. While SocEvac may a bit ahead of its time, the data will become available to do more frequent and larger scale studies in the near future.

6.4 Conclusion

This model, while still in the prototype stage, allowed testing of the hypotheses developed in this dissertation, and makes a strong case for the incorporation of collective behavior, groups, and supra force in models of building evacuation. As hypothesized, in general members of groups take more time to exit buildings than unaffiliated individuals. Also, the social behavior model leads to improved aggregate results of exit use compared to closest exit platforms and most other available modeling platforms of the Station.

Social behavior modeling also does a better job of explaining the conditions that occurred, at least during the Station scenario, that caused so many occupants to use exits that were not close to their initial location. As a result of this, the SocEvac model did a better job of explaining individual exit paths than conventional models. Although this was not a profound difference, it is a step in the right direction for building evacuation modeling.

Today, the SocEvac model is ready for other scenarios, and researchers will need to find other data sources with as much detail as the Station evacuation. While these data are not available today, location data collected from networked devices will provide ample opportunities for event recreation. In the mean time, the SocEvac code requires numerous refinements, and there is much to be done to change the definition of "social behavior" in building evacuation modeling.

REFERENCES

Abdelhak, Haifa, Eric Daude, and Damien Olivier. 2009. “Exploring Crowd Dynamics Based on Individual Behaviors and Local Interactions.” Le Havre, France: University of Le Havre.

Aguirre, B. E. 2005. "Emergency evacuations, panic, and social psychology." Psychiatry-Interpersonal and Biological Processes no. 68 (2):121-129.

Aguirre, B. E., Sherif El-Tawil, Eric Best, K. B. Gill, and Vladamir Federov. 2011. "Contributions of Social Science to Agent-Based Models of Building Evacuation." Contemporary Social Science no. 6 (3):18.

Aguirre, B. E., M. Torres, and K. B. Gill. 2009. “A Test of Pro Social Explanation of Human Behavior in Building Fire.” Paper read at NSF Engineering Research and Innovation Conference.

Aguirre, B. E., Manuel R. Torres, Kimberly B. Gill, and H. Lawrence Hotchkiss. 2011. "Normative Collective Behavior in The Station Building Fire." Social Science Quarterly no. 92 (1):100-118.

Al-Obaidi, S., J. C. Wall, A. Al-Yaqoub, and M. Al-Ghanim. 2003. "Basic gait parameters: A comparison of reference data for normal subjects 20 to 29 years of age from Kuwait and Scandinavia." Journal of Rehabilitation Research and Development no. 40 (4).

Apple Corporation. 2011. "iOS 4.3.3 Software Update." List of improvements from iOS software version 4.3.2. United States: Apple Inc.

ARA. 2012. E-Sim Event Simulator. In http://www.ara.com/products/e-sim/esim.pdf, edited by Applied Research Associates.

Argonne. 2012. Repast. Argonne National Laboratory 2011 [cited January 29 2012]. Available from http://repast.sourceforge.net/.

AT&T. 2012. "AT&T Location Information Services." Description of location information services from AT&T. United States: AT&T.

Bacharach, Michael. 2006. Beyond individual choice: teams and frames in game theory. Princeton, NJ: Princeton University Press.

Balci, O. 1997. "Principles of simulation model validation, verification, and testing." Transactions of the Society for Computer Simulation no. 14 (1):3-12.

Ball, Richard, and Kateryna Chernova. 2008. "Absolute income, relative income, and happiness." Social Indicators Research no. 88 (3):497-529.

Banks, Jerry, John Carson, Barry L. Nelson, and David Nicol. 2004. Discrete-Event System Simulation. Edited by W. J. Fabrycky and J. H. Mize. 4th Edition ed, Prentice Hall International Series in Industrial and Systems Engineering. Upper Saddle River, New Jersey: Prentice Hall.

Barnshaw, John, Lynn Letukas, and E.L. Quarantelli. 2008. The Characteristics of Catastrophes and Their Social Evolution: An Exploratory Analysis of Implications for Crisis Policies and Emergency Management Procedures. In Disaster Research Center Working Paper Series. Newark, DE: University of Delaware.

Barylick, John. 2012. Killer Show. Lebanon, NH: University Press of New England.

Beaulieu, C. M. J. 2004. "Intercultural study of personal space: A case study." Journal of Applied Social Psychology no. 34 (4).

Best, Eric. 2012. Collective Models of Disaster: Making a Case for Using Collective Mobile Phone Location Data in Disaster Science. http://www.colorado.edu/hazards/awards/paper-competition/best.pdf

Binmore, Ken. 2007. Does game theory work? The bargaining challenge. Cambridge, MA: The MIT Press.

Bohannon, J. 2005. "Building safety - Directing the herd: Crowds and the science of evacuation." Science no. 310 (5746):219

Bonabeau, E. 2002. "Agent-based modeling: Methods and techniques for simulating human systems." Proceedings of the National Academy of Sciences of the United States of America no. 99:7280-7287.

Bonato, Anthony, and Richard J. Nowakowski. 2011. The game of cops and robbers on graphs. Providence, RI: American Mathematica Society.

Brogan, D. C., and J. K. Hodgins. 1997. "Group behaviors for systems with significant dynamics." Autonomous Robots no. 4 (1):137-153.

CADDPrimer. 2012. CADD Primer. Mailmax Publishing 2011 [cited January 29 2012]. Available from http://www.caddprimer.com/cadd_primer_chapters/drawing_design_automation_introduction_0.htm.

Calabrese, Francesco, Zbigniew Smoreda, Vincent D. Blondel, and Carlo Ratti. 2011b. "Interplay between Telecommunications and Face-to-Face Interactions: A Study Using Mobile Phone Data." *Plos One* 6.

Cardona, Omar D., Mario G. Ordaz, Luis E. Yamin, Mabel C. Marulanda, and Alex H. Barbat. 2008. "Earthquake loss assessment for integrated disaster risk management." Journal of Earthquake Engineering no. 12:48-59.

Castiglione, F. 2011. Agent Based Modeling. Scholarpedia. Sponsored by Eugene M. Izhikevich. 2010 [cited December 6 2011].

Chaturvedi, A., A. Mellema, and S. Filatyev. 2006. "DDDAS for Fire and Agent Evacuation Modeling of the Rhode Island Nightclub Fire." ICCS 2006, Part III, LNCS 3993. Berlin: Springer-Verlag, p.433-439.

Chen, C. M., and Y. Yu. 2000. "Empirical studies of information visualization: a meta-analysis." International Journal of Human-Computer Studies no. 53 (5):851-866.

Cheng, H., and G. V. Hadjisophocleous. 2011. "Dynamic modeling of fire spread in building." Fire Safety Journal no. 46 (4):211-224.

Cherif, Foudil, and Noureddine Djedi. 2005. A Framework to Simulate the Evacuation of a Crowd in Emergency Situations.

Chu, YanYan, Hui Zhang, ShiFei Shen, Rui Yang, and LiFeng Qiao. 2010. "Development of a model to generate a risk map in a building fire." Science China-Technological Sciences no. 53 (10):2739-2747.

Colet, E., and D. Aaronson. 1995. "Visualization of Multivariate Data – Human-Factors Considerations." Behavior Research Methods Instruments & Computers no. 27 (2):257-263.

Cooper, E. W., and K. Kamei. 2002. "A study of color conspicuity for ease-of-use inference in visualization." Color Research and Application no. 27 (2):74-82.

Davis, John B. 2011. Individuals and identity in economics. New York, NY: Cambridge University Press.

de Silva, F. N., and R. W. Eglese. 2000. "Integrating simulation modelling and GIS: spatial decision support systems for evacuation planning." Journal of the Operational Research Society no. 51 (4):423-430.

DeVito, Joseph A. 2003. The Interpersonal Communication Book. Tenth Edition ed. Boston, MA: Allyn & Bacon.

dos Santos, S., and K. Brodlie. 2004. "Gaining understanding of multivariate and multidimensional data through visualization." Computers & Graphics-Uk no. 28 (3):311-325.

Drevermann, H., and D. Travis. 1998. "Visualization using colour: visual presentation of events in particle physics." Behaviour & Information Technology no. 17 (1):18-26.

Drury, John, Chris Cocking, and Steve Reicher. 2009. "Everyone for themselves? A comparative study of crowd solidarity among emergency survivors." British Journal of Social Psychology no. 48 (3):487-506.

Duval, Robert F. 2006. NFPA Case Study: Nightclub Fires. National Fire Protection Association.

Ehibara, Manabu, Akira Ohtsuki, and Hideaki Iwaki. 1992. "A Model for Simulating Human Behavior During Emergency Evacuation Based on Classificatory Reasoning and Certainty Value Handling." Computer-Aided Civil and Infrastructure Engineering no. 7 (1):8.

Entertainment, Blizzard. 2011. World of Warcraft. In Warcraft: Blizzard Entertainment.

Ewer, John, Edwin R. Galea, Fuchen Jia, Peter Lawrence, Ana Veeraswamy, and Zhaozhi Wang. 2008. "Coupled Fire/Evacuation Analysis of the Station Nightclub Fire." Fire Safety Science no. 9:465-476.

Fahy, R.F., G. Proulx, and J. Flynn. 2011. "The Station Nightclub Fire - An Analysis of Witness Statements." Fire Safety Science no. 10:197-209.

Fang, Zhi-Ming, Wei-Guo Song, Zi-Juan Li, Wei Tian, Wei Lv, Jian Ma, and Xia Xiao. 2012. "Experimental study on evacuation process in a stairwell of a high-rise building." Building and Environment no. 47:316-321.

Fang, Zhiming, Weiguo Song, Jun Zhang, and Hao Wu. 2010. "Experiment and modeling of exit-selecting behaviors during a building evacuation." Physica a-Statistical Mechanics and Its Applications no. 389 (4):815-824.

Fruin, J. 1987. Pedestrian Planning and Design. Mobile, AL: Elevator World.

Galea, E. R. 2004. Computational Fire Engineering - Do We Have What We Need? In 3rd International Symposium on Human Behavior in Fire. Belfast, UK.

Galea, E. R., G. Sharp, and P. J. Lawrence. 2008. "Investigating the Representation of Merging Behavior at the Floor-Stair Interface in Computer Simulations of Multi-Floor Building Evacuations." Journal of Fire Protection Engineering no. 18 (4):291-316.

Galea, E.R., Z. Wang, and A. Veeraswamy. 2008. “Coupled Fire/Evacuation Analysis of the Station Nightclub Fire.” Fire Safety Science 2008(9):465-476.

Gill, Kimberly B., Elizabeth A. Laposata, Christina F. Dalton, and Benigno E. Aguirre. 2010. Analysis of Severe Injury and Fatality in The Station Nightclub Fire. Newark, DE: Univeristy of Delaware.

Grosshandler, W., N. Bryner, D. Madrizykowski, and K. Kuntz. 2005. A Report of the Technical Investigation of the Station Nightclub Fire. In NCSTAR2(1): NIST.

Guo, Diansheng. 2009. "Flow Mapping and Multivariate Visualization of Large Spatial Interaction Data." Ieee Transactions on Visualization and Computer Graphics no. 15 (6):1041-1048.

Gwynne, S., E. R. Galea, and P. J. Lawrence. 2006. "The introduction of social adaptation within evacuation modelling." Fire and Materials no. 30 (4):285-309.

Gwynne, S., E. R. Galea, P. J. Lawrence, and L. Filippidis. 2001. "Modelling occupant interaction with fire conditions using the buildingEXODUS evacuation model." Fire Safety Journal no. 36 (4):327-357.

Gwynne, S., E.R. Galea, P. Lawrence, and L. Filippidis. 2000. Modeling Occupant Interaction with Fire Conditions Using the buildingEXODUS Evacuation Model.

Ha, Vi, and George Lykotrafitis. 2012. "Agent-based modeling of a multi-room multi-floor building emergency evacuation." Physica a-Statistical Mechanics and Its Applications no. 391 (8):2740-2751.

Haron, Fazilah, Yasser M. Alginahi, Muhammad N. Kabir, and Ali I. Mohammed. 2012. "Software Evaluation for Crowd Evacuation - Case Study: Al-Masijd An-Nabawi." International Journal of Computer Science Issues no. 9 (6):7.

Hart, P.E., N.J. Nilsson, and B. Raphael. 1968. "A Formal Basis for the Heuristic Determination of Minimum Cost Paths." IEEE Transactions on Systems Science and Cybernetics SSC4 no. 4 (2):8.

Helbing, D., I. Farkas, and T. Vicsek. 2000. "Simulating dynamical features of escape panic." Nature no. 407 (6803):487-490.

Helbing, Dirk, Anders Johansson, and Habib Zein Al-Abideen. 2007. "Dynamics of crowd disasters: An empirical study." Physical Review E no. 75 (4).

Henderson, L.F. 1971. "Statistics of Crowd Fluids." Nature no. 229 (5284):381-&.

Isobe, M., D. Helbing, and T. Nagatani. 2004. "Experiment, theory, and simulation of the evacuation of a room without visibility." Physical Review E no. 69 (6).

Jehiel, Philippe, and Frederic Koessler. 2008. "Revisiting games of incomplete information with analogy-based expectations." Games and Economic Behavior no. 62 (2):533-557.

Johnson, N. R., and W. E. Feinberg. 1997. "The impact of exit instructions and number of exits in fire emergencies: A computer simulation investigation." Journal of Environmental Psychology no. 17 (2):123-133.

Johnson, Norris R., W. E. Feinberg, and D. M. Johnston. 1994. "Microstructure and Panic: The Impact of Social Bonds on Individual Action in Collective Flight from the Beverly Hills Supper Club Fire." In Disasters, Collective Behavior, and Social Organization, edited by R. R. Dynes and K. Tierney, 168 - 189. Newark, Delaware: University of Delaware Press.

Kent, R. L., and S. E. Moss. 1994. "Effects of Sex and Gender-Role on Leader Emergence." Academy of Management Journal no. 37 (5):1335-1346.

Kleijnen, J. P. C. 1995. "Verification and Validation of Simulation-Models." European Journal of Operational Research no. 82 (1):145-162.

Kukla, R., J. Kerridge, A. Willis, J. Hine, and Trb. 2001. "PEDFLOW - Development of an autonomous agent model of pedestrian flow." Artificial Intelligence and Intelligent Transportation Systems: Planning and Administration (1774):11-17.

Kuligowski, Erica D., and Steve M.V. Gwynne. 2010. "The Need for Behavioral Theory in Evacuation Modeling." In Pedestrian and Evacuation Dynamics 2008, edited by W.W.F. Klingsch. Springer.

Kuligowski, Erica D., and Richard D. Peacock. 2005. A Review of Building Evacuation Models. In Technical Note, edited by National Institute of Standards and Technology. Washington, DC: National Institute of Standards and Technology.

Lee, Teng-Yok, and Han-Wei Shen. 2009. "Visualization and Exploration of Temporal Trend Relationships in Multivariate Time-Varying Data." Ieee Transactions on Visualization and Computer Graphics no. 15 (6):1359-1366.

Li, Huibo, Wen Tang, and Derek Simpson. 2004. Behavior Based Motion Simulation for Fire Evacuation Procedures. Paper read at Theory and Practice of Computer Graphics.

Lord, J., B.J. Meacham, A. Moore, R.F. Fahy, and G. Proulx. 2005. Guide for Evaluating the Predictive Capabilities of Computer Egress Models. Washington, DC: National Institute of Standards and Technology.

Low, D. J. 2000. "Statistical physics - Following the crowd." Nature no. 407 (6803):465-466.

Luo, L. B., S. P. Zhou, W. T. Cai, M. Y. H. Low, F. Tian, Y. W. Wang, X. Xiao, and D. Chen. 2008. "Agent-based human behavior modeling for crowd simulation." Computer Animation and Virtual Worlds no. 19 (3-4):271-281.

Manocha, Dinesh, and Ming C. Lin. 2012. "Interaction Large-Scale Crowd Simulation." Digital Urban Modeling and Simulation - Communications in Computer and Information Science no. 242:15.

Mawson, A. R. 2005. "Understanding mass panic and other collective responses to threat and disaster." Psychiatry-Interpersonal and Biological Processes no. 68 (2):95-113.

McCauley Bush, P., M. Bull, A. Soyler, Y. Zhu, and S.Kanjanapongpaisal Diaz, P. 2009. "A Research-Based Approach to Predictive Simulation in Disaster Management." Soft Computing Applications no. 2009:2.

Meacham, B., J. Lord, A. Moore, G. Proulx, R. Fahy, and K. Notarianni. 2004. "Investigation of Uncertainty in Egress Models and Data." Proceedings of the 3rd International Symposium on Human Behaviour in Fire. Belfast, UK. pp.419-428.

Mihram, G. A. 1972. "Some Practical Aspects of Verification and Validation of Simulation Models." Operational Research Quarterly no. 23 (1):17

Miltiadis, A. B. 2010. "A 3D Fire Spread Model Implementing Vegetation Combustion applicable to Mountainous Forest Fires." Disaster Advances no. 3 (3):54-63.

Mott MacDonald. 2012. STEPS Software - Simulating Pedestrian Dynamics.

Moussaid, Mehdi, Niriaska Perozo, Simon Garnier, Dirk Helbing, and Guy Theraulaz. 2010. "The Walking Behaviour of Pedestrian Social Groups and Its Impact on Crowd Dynamics." Plos One no. 5 (3).

Murphy, Alan, and KJ Salchow Jr. 2007. “Applied Application Security-- Positive and Negative Efficiency.” White Paper. United States: F5 Networks, Inc.

Nielsen, Helle. 2009. Bonded Rationality in Decision-Making. Manchester, UK: Manchester University Press.

Notarianni, Kathy A. 2000. “The Role of Uncertainty in Improving Fire Protection Regulation.” Dissertation. Engineering and Public Policy, Carnegie Mellon University, Pittsburgh, PA.

Oberg, T., A. Karsznia, and K. Oberg. 1993. “Basic Gait Parameters – Reference Data for Normal Subjects, 10-79 Years of Age." Journal of Rehabilitation Research and Development no. 30 (2).

Olsson, P. A., and M. A. Regan. 2001. "A comparison between actual and predicted evacuation times." Safety Science no. 38 (2):139-145.

Owen, Matthew, Edwin R. Galea, and Peter J. Lawrence. 1996. "The Exodus Evacuation Model Applied To Building Evacuation Scenarios." Journal of Fire Protection Engineering no. 8 (2):19.

Pan, X. 2010. "Walk This Way." Passenger Terminal World no. 2010:6.

Pan, X. H., C. S. Han, K. Dauber, and K. H. Law. 2006. "Human and social behavior in computational modeling and analysis of egress." Automation in Construction no. 15 (4):448-461.

Pan, Xiaoshan, Charles S. Han, Ken Dauber, and Kincho H. Law. 2007. "A multi-agent based framework for the simulationof human and social behaviors during emergency
evacuations." AI & Soc no. 22:113-132.

Park, Y. J., A. H. S. Ang, and Y. K. Wen. 1985. "Seismic Damage Analysis of Reinforced-Concrete Buildings." Journal of Structural Engineering-Asce no. 111 (4):740-757.

Pelechano, Nuria, and Norman I. Badler. 2006. "Modeling crowd and trained leader behavior during building evacuation." Ieee Computer Graphics and Applications no. 26 (6):80-86.

Pelechano, Nuria, and Ali Malkawi. 2008. "Evacuation simulation models: Challenges in modeling high rise building evacuation with cellular automata approaches." Automation in Construction no. 17 (4):377-385.

Pelechano, Nuria, Kevin O'Brien, Barry Silverman, and Norman Badler. 2005. Crowd Simulation Incorporating Agent Psychological Models, Roles and Communication. University of Pennsylvania: Center for Human Modeling and Simulation.

Pham, Tuan, Rob Hess, Crystal Ju, Eugene Zhang, and Ronald Metoyer. 2010. "Visualization of Diversity in Large Multivariate Data Sets." Ieee Transactions on Visualization and Computer Graphics no. 16 (6):1053-1062.

Pidd, Michael. 2010. Modelling in Management Science. Third Edition ed. Hoboken, NJ: Wiley.

Pimenta, Luciano C. A., Nathan Michael, Renado C. Mesquita, Guilherme A. S. Pereira, and Vijay Kumar. 2008. Control of Swarms Based on Hydrodynamic Models. In IEEE International Conference of Robotics and Automation. Pasadena, CA: IEEE.

Pinelli, J. P., E. Simiu, K. Gurley, C. Subramanian, L. Zhang, A. Cope, J. J. Filliben, and S. Hamid. 2004. "Hurricane damage prediction model for residential structures." Journal of Structural Engineering-Asce no. 130 (11):1685-1691. doi: 10.1061/(asce)0733-9445(2004)130:11(1685).

Pineo, Daniel, and Colin Ware. 2012. "Data Visualization Optimization via Computational Modeling of Perception." Ieee Transactions on Visualization and Computer Graphics no. 18 (2):309-320.

Predtechenskii, V.M., and A.I. Milinskii. 1978. Planning for Foot Traffic Flow in Buildings. New Delhi: Amerind.

PTV. 2012. VISSIM. PTV AG: PTV Vision 2012 [cited January 13 2012]. Available from http://www.vissim.de/software/transportation-planning-traffic-engineering/software-system-solutions/vissim/.

Quarantelli, E. L. 1960. "IMAGES OF WITHDRAWAL BEHAVIOR IN DISASTERS - SOME BASIC MISCONCEPTIONS." Social Problems no. 8 (1):68-79.

Quarantelli, E.L. 2008. "Conventional Beliefs and Counterintuitive Realities." Social Research no. 75 (3):873-904.

Ren, Chuanjun, Chenghui Yang, and Shiyao Jin. 2009. "Agent-Based Modeling and Simulation on Emergency Evacuation." Complex no. 5:11.

Riley, P. O., U. Della Croce, and D. C. Kerrigan. 2001. "Propulsive adaptation to changing gait speed." Journal of Biomechanics no. 34 (2).

Robinson, S., and R. J. Brooks. 2010. "Independent Verification and Validation of an Industrial Simulation Model." Simulation-Transactions of the Society for Modeling and Simulation International no. 86 (7):405-416.

Rucker, R. 2004. Infinity and the Mind: The Science and Philosophy of the Infinite. Princeton, NJ: Princeton Science Library.

Sagun, Aysu, Dino Bouchlaghem, and Chimay J. Anumba. 2011. "Computer simulations vs. building guidance to enhance evacuation performance of buildings during emergency events." Simulation Modelling Practice and Theory no. 19 (3):1007-1019.

Santos, Gabriel, and Benigno E. Aguirre. 2004. A Critical Review of Emergency Evacuation Simulation Models. In DRC Preliminary Papers. Newark, DR: Disaster Research Center, University of Delaware.

Sargent, Robert. 2008. "Verification and Validation of Simulation Models." In Winter Simulation Conference, edited by R. R. Hill S. J. Mason, L. Mönch, O. Rose, T. Jefferson, J. W. Fowler. Miami, FL: IEEE.

Schadschneider, Andreas, Wolfram Klingsch, Hubert Klupfel, Tobias Kretz, Christian Rogsch, and Armin Seyfried. 2008. "Evacuation Dynamics: Empirical Results, Modeling and Applications." In Encylopedia of Complexity and System Science, edited by B. Meyers. Berlin, Germany: Springer.

Sengupta, Pratim. 2011. Design Principles for a Visual Programming Language to Integrate Agent-Based Modeling in K-12 Science. edited by Matter & Media Lab Mind. Nashville, TN: Vanderbilt University.

Shih, N. J., C. Y. Lin, and C. H. Yang. 2000. "A virtual-reality-based feasibility study of evacuation time compared to the traditional calculation method." Fire Safety Journal no. 34 (4):377-391.

Silva, Samuel, Beatriz Sousa Santos, and Joaquim Madeira. 2011. "Using color in visualization: A survey." Computers & Graphics-Uk no. 35 (2):320-333.

Silveira, R., E. Prestes, and L. P. Nedel. 2008. "Managing coherent groups." Computer Animation and Virtual Worlds no. 19 (3-4):295-305.

Simon, Herbert A. 1955. "A Behavioral Model of Rational Choice." Quarterly Journal of Economics no. 69 (1):99-118.

Sommer, R. 1959. "Studies in Personal-Space." Sociometry no. 22 (3).

Song, W. G., X. Xu, B. H. Wang, and S. J. Ni. 2006. "Simulation of evacuation processes using a multi-grid model for pedestrian dynamics." Physica a-Statistical Mechanics and Its Applications no. 363 (2):492-500.

Spearpoint, M. J. 2012. "Network modeling of The Station Nightclub fire evacuation." Journal of Fire Protection Engineering no. 22 (3).

Stirling, M. W., G. H. McVerry, and K. R. Berryman. 2002. "A new seismic hazard model for new Zealand." Bulletin of the Seismological Society of America no. 92 (5):1878-1903.

Swarm. 2012. Swarm Software for Agent-based Modeling. Center for the Study of Complex Systems, University of Michigan 2011 [cited January 29 2012]. Available from http://www.swarm.org/index.php/Main_Page.

Thelwall, Mike, David Wilkinson, and Sukhvinder Uppal. 2010. "Data Mining Emotion in Social Network Communication: Gender Differences in MySpace." *Journal of the American Society for Information Science and Technology* 61:190-199.

Thunderhead Engineering. 2012. Verification and Validation. Pathfinder 2012.1: 33-36

Tissera, Pablo Christian, A. Marcela Printista, and Emilio Luque. 2012. "A Hybrid Simulation Model to Test Behaviour Designs in an Emergency Evacuation." Procedia Computer Science no. 9 (2012):10.

Torres, Manuel R. 2010. Every Man for Himself? Testing Multiple Conceptual Approaches of Emergency Egress on Building Evacuation During a Fire, Department of Sociology, University of Delaware, Newark, DE.

Tory, M., and T. Moller. 2004. "Human factors in visualization research." Ieee Transactions on Visualization and Computer Graphics no. 10 (1):72-84.

Tsai, Jason, Natalie Friedman, Emma Bowring, Matthew Brown, Shira Epstein, Gal Kaminka, Stacy Marsella, Andrew Ogden, Inbal Rika, Ankur Sheel, Matthew Taylor, Xuezhi Wang, Avishay Zilka, and Milind Tambe. 2011. ESCAPES - Evacuation Simulation with Children, Authorities, Parents, Emotions, and Social Comparison. Paper read at 10th International Conference on Autonomous Agents and Multiagent Systems - Innovative Applications Track, at Taipei, Taiwan.

Tubbs, Jeffrey S. and Brian J. Meacham. 2007. “Egress Design Solutions.” Hoboken, NJ: Wiley.

Tueretgen, Ilknur Oezalp, Pinar Unsal, and Inci Erdem. 2008. "The effects of sex, gender role, and personality traits on leader emergence - Does culture make a difference?" Small Group Research no. 39 (5):588-615.

Valckenaers, P., J. Sauter, C. Sierra, and J. A. Rodriguez-Aguilar. 2007. "Applications and environments for multi-agent systems." Autonomous Agents and Multi-Agent Systems no. 14 (1):61-85.

Wang, Yanhui, Liqiang Zhang, Jingtao Ma, Liu Liu, Lixin Zhang, and Dongqin You. 2011. "Combining Building and Behavior Models for Evacuation Planning." Ieee Computer Graphics and Applications no. 31 (3):42-55.

Wann, J., and M. MonWilliams. 1996. "What does virtual reality NEED?: Human factors issues in the design of three-dimensional computer environments." International Journal of Human-Computer Studies no. 44 (6):829-847.

Warden, Pete. 2011. "iPhone Tracker." Instructions to obtain basic location tracking database from iOS 4.3.2 software. United Kingdom: Github.

Webber, G.M.B., M.S. Wright, and G.K. Cook. 2001. Effects of Smoke on People's Walking Speeds Using Overhead Lighting and Wayguidance Provision. In Human Behavior in Fire: Proceedings of the 2nd International Symposium. London: Interscience Communications.

Weller, J. M., and E. L. Quarantelli. 1973. "Neglected Characteristics of Collective Behavior." American Journal of Sociology no. 79 (3):665-685.

Wilensky, Uri. 2012. NetLogo. Northwestern University Center for Connected Learning and Computer-Based Modeling 2011 [cited January 29 2012]. Available from http://ccl.northwestern.edu/netlogo/.

Xu, Xuan, and Weiguo Song. 2009. "Staircase evacuation modeling and its comparison with an egress drill." Building and Environment no. 44 (5):1039-1046.

Yang, L. Z., D. L. Zhao, J. Li, W. F. Fang, and W. C. Fan. 2004. "Simulation of evacuation behaviors in fire using spacial grid." Progress in Natural Science no. 14 (7):614-618.

Youngblood, S. M., and D. K. Pace. 1995. "An Overview of Model and Simulation Verification, Validation, and Accreditation." Johns Hopkins Apl Technical Digest no. 16 (2):197-206.

Zhang, Huiqi, Ram Dantu, and Joao W. Cangussu. 2011. "Socioscope: Human Relationship and Behavior Analysis in Social Networks." *Ieee Transactions on Systems Man and Cybernetics Part a-Systems and Humans* 41:1122-1143.

Zheng, Xiaoping, Tingkuan Zhong, and Mengting Liu. 2009. "Modeling crowd evacuation of a building based on seven methodological approaches." Building and Environment no. 44 (3):437-445.