Agent-based Evacuation Modeling: Simulating the Los Angeles International Airport

Jason Tsai, Emma Bowring, Shira Epstein, Natalie Fridman, Prakhar Garg, Gal Kaminka, Andrew Ogden, Milind Tambe, Matthew Taylor

{jasontts@usc.edu, ebowring@pacific.edu, spepstei@usc.edu, fridman@cs.biu.ac.il, prakhargarg.jiit@gmail.com, galk@cs.biu.ac.il, aogden@usc.edu, tambe@usc.edu, taylorm@usc.edu}

Introduction

In the aftermath of a terrorist attack on a large public venue such as an airport, a train station, or a theme park, rapid but safe evacuation is critical. For example, multiple IED explosions at the Los Angeles International Airport (LAX) will require evacuation of thousands of travelers and airport staff, while mitigating the risks from possible secondary explosions. In such cases, it is crucial to obtain information on the fastest evacuation routes, the time needed to evacuate people, the lag between the disaster and the arrival of a bomb squad or other emergency personnel, and what tactics and procedures authorities can use to avoid stampedes, confusion, and loss of life. By understanding possible scenarios in simulation before hand, emergency personnel can be trained so that they respond in the event of an actual evacuation.

Unfortunately, conducting live exercises for evacuating thousands of people is generally impossible. It would be time-consuming and would result in tremendous lost revenue. Moreover, a staged evacuation would necessarily miss crucial aspects of the required response (e.g. fear, confusion) on the part of both emergency personnel and crowd evacuees, or the exercise would be considered unethical. Smaller scale evacuation exercises miss the most important feature of an evacuation: its massive scale.

Simulations provide one attractive answer. Evacuation simulations allow us to meet the required training, evacuation planning, and tactics development goals; by running large numbers of “faster than real-time” simulations, we can obtain data from large numbers of scenarios that would be near-impossible in live exercises. This information can be used by policy-makers to predetermine optimal solutions to time-critical situations such as those involving injuries, IEDs, or active shooters. Additionally, real-time simulations can be created for officers on the ground who may only see a handful of real emergencies in their careers and thus could benefit immensely from repeated scenario-style training tools to learn with. In building these simulations, we bring to bear over two-decades of experience in agent-based simulations, including battlefield simulations of battalions of virtual helicopter pilots or teams of virtual fighter pilots for DARPA’s Synthetic Theater of War program and disaster response simulations in collaboration with the Los Angeles Fire Department [7].

Previous Work

Current state-of-the-art crowd simulations, despite useful advances, suffer from many drawbacks. Broadly speaking, existing work can be classified as “macro-oriented” or “micro-oriented.”

Macro-oriented simulations have seen much more attention in the past, likely due to computational constraints that previously prohibited a fine-grained treatment of simulations. This approach has been very successfully used in leading evacuation simulations and generally allow for control of all entities in the simulation as a whole, but do not accommodate for detailed simulation of individual intelligent agents [4]. For many applications, such as egress modeling for architectural planning, this is precisely what is desired --- the time-to-evacuate metric is both directly relevant and can be accurately measured with this type of crowd treatment.

When applied to training and policy-making, however, several key shortcomings become apparent. Abstract simulations can be challenging to interpret from a policy-making standpoint and can be difficult for officers to train with and condition on. In a training simulation, visual realism becomes highly desirable, as officers who have seen a situation before, even virtually, will be better prepared to respond. Furthermore, macro-level simulations assume a great deal about the capabilities of an individual agent, such as that every agent always knows where the nearest exit is and proceeds to it without hesitation. However, in real life, the majority of an officer's duty is to direct confused and lost people to the most appropriate exit, which is not always the nearest one. Not capturing this aspect of human behavior makes the simulation inaccurate and less useful for training purposes. Many other micro-level details also come into play that may alter the situation from an officer-on-the-ground's perspective. Some examples are: reasoning by sub-groupings like families or friends, the impact of heightened emotions on locomotion (e.g., pushing and other competitive crowd behavior), the response to authority figures, the impact of baggage, and cognitive shortcomings with respect to path-planning and identifying signage. In addition, even some crucial macro-level behaviors are not captured in existing simulators. For example, neither modeling regular activities *prior* to an emergency nor understanding how people transition from obliviousness to understanding that they are in an emergency situation are included in the current generation of simulators.

There has also been substantial work in micro-oriented approaches, many of which originated in the graphics community. Perhaps the most publicized has been Massive, the company behind the vast majority of the cinema world's crowd scenes, including many evacuation scenarios. Their simulations are seen in movies such as "Lord of the rings", "I, Robot", "Die Hard 4.0", and "Wall-E." In the graphics community, the goal is dramatic appeal and believability, for which Massive works splendidly, as evidenced by their commercial success and their customers' box office numbers. However, for training and policy-making, realism takes center-stage. Simply because a digital character's actions during an emergency *look* real is no longer sufficient validation. As stated before, an officer must train on something that exhibits real-life behavior, not only plausible or dramatic behavior. Similarly, a policy-maker making decisions on inaccurate models may lead to errors that cost hundreds of human lives.

Research Approach

A crucial aspect of our research methodology is "use-inspired research". Starting with a concrete real-world problem focuses our research to problems of importance in the real-world, while simultaneously providing us avenues for obtaining real-world input data, and potentially real-world validation of our results. To that end, we have established collaboration with the Los Angeles International Airport, which began with the implementation of the ARMOR security assistant for randomization of personnel placement [6]. We have already had several discussions with officials of the Transportation Security Administration at LAX and with the Los Angeles World Airports police department, two primary agencies dealing with evacuation modeling at LAX. This will allow us detailed access to data at LAX as well as current plans for airport evacuation at LAX.

Micro-oriented, agent-based simulations enable the representation of individual idiosyncrasies and realistically model evacuation scenarios. Endowing each agent with different behavioral tendencies provides an initial step towards overcoming the micro-level difficulties and has seen some traction in existing literature. We will use a BDI style agent to achieve realistic, individual behaviors and allow for heterogeneity by providing sets of responses for which each agent can have differing tendencies to follow. Coupled with each agent's own (limited) perception abilities, our simulation will have a very wide range of realistic agent behaviors.

The next step will be to incorporate and understand the impact on small-scale social interactions, which has not been explored and plays a key role in how airport evacuations unfold. For example, during an emergency, observing others rushing to the exits will affect the observer's decision-making process and will probably cause him to do the same. This contagious behavior was found to be one of the keystones of crowd behavior [5] [1]. To provide our agents with social skills crucial to an accurate simulation of evacuation behavior, we utilize a computational model of social comparison processes (SCT). The SCT model [3] is based on Social Comparison Theory, a popular social psychology theory that was initially presented by Festinger in 1954 [2] and has been continuously evolving since then. The key idea is that humans, lacking objective means to evaluate their state, compare themselves to others that are similar to them, and then take actions to reduce the differences found. We believe that social comparison is a general cognitive process underlying social behavior. Thus, we can use SCT to model the characteristics of crowd behavior such as contagious behavior and the spread of awareness when an emergency first arises as well as variations in sub-groupings such as friends and families.

Also, we will collaborate with Massive Software to produce realistic graphics for the front-end of our simulations. We recognize the importance of visual integrity in training an officer and policy-making and use Massive's Software's extensive capabilities to couple realistic behaviors with realistic graphics and produce a new standard for simulations.

Finally, we intend on validating our simulation as rigorously as possible. Short of real-life evacuation data which is not available due to security concerns, we must rely on secondhand knowledge and comparison techniques. One form of this is using the aforementioned macro-oriented simulations. Despite their weaknesses, they are very accurate in modeling the high-level behavior of crowds. Thus, any simulation we create must not only mimic the individual behavior of agents, but must also match the metrics measured by these macro-oriented simulations. We intend on implementing the state-of-the-art in existing macro-level simulations as a benchmark for our own work as part of our validation process. Furthermore, we intend on leveraging airport security officials' extensive experience with evacuation scenarios to evaluate and rate our simulations.

Anticipated Challenges

We foresee three major challenges which are not unique to our approach, but are made difficult by the domain itself. The first is computational speed. Past work has cited this as an issue when reaching 20 or 30 agents due to the complexity of an involved BDI-architecture. Any given LAX airport terminal has hundreds if not thousands of agents and our approach will use an even more complex decision mechanism than past approaches. Thus we will need to develop more efficient computational methods or streamline the decision mechanism in some other way for our simulation to be practically useful.

Second, the calibration of parameters will be difficult to do in a scientifically-grounded way. While all simulations suffer from this problem, our domain, where raw video data of relevant situations is confidential for security reasons, makes this significantly more difficult. Thus, instead of analyzing actual footage of these types of situations to evaluate human behavior, we will need to use a combination of domain expert advice and quantitative psychological studies of evacuations and Social Comparison.

Finally, validation of our simulation will also be made difficult by the nature of our domain. By using domain experts to help with calibration, validation for actual users and security experts should be more convincing. Ideally, we would accompany this with results from controlled experiments of the situation being simulated that are quantitatively compared with simulation predictions. However, additional validation with experimentation will be difficult to achieve, since experiments of evacuation situations will necessarily lack realism if people realize it is a controlled experiment or be unethical to perform if they are not informed.

References:

[1] F. H. Allport. *Social Psychology*. Boston: Houghton Mifflin, 1924.

[2] L. Festinger. A theory of social comparison processes. *Human Relations*, 117-140, 1954.

[3] N. Fridman and G. Kaminka. Towards a Cognitive Model of Crowd Behavior Based on Social Comparison Theory. In *AAAI-07*, 2007.

[4] D. Helbing, I. Farkas, and T. Vicsek. Simulating Dynamical Features of Escape Panic. *Nature*. 407:487-490, Sept. 2000.

[5] G. Le Bon. *The crowd: A study of the popular mind.* Dunwoody, Ga., N.S. Berg, 1895.

[6] J. Pita, M. Jain, C. Western, C. Portway, M. Tambe, F. Ordonez, S. Kraus, and P. Parachuri. Depoloyed ARMOR protection: The application of a game-theoretic model for security at the Los Angeles International Airport. In *AAMAS-08 (Industry Track)*, 2008.

[7] N. Schurr, P. Patil, F. Pighin, M. Tambe. Using multiagent teams to improve the training of incident commanders. *In AAMAS-06*, 2006.