Proceedings of the 2018 Winter Simulation Conference
M. Rabe, A.A. Juan, N. Mustafee, A. Skoogh, S. Jain, and B. Johansson, eds.

AGENT-BASED MODELING OF A STADIUM EVACUATION IN A SMART CITY

Wenxin Zhang
Vincent Terrier
Xiaoshu Fei
Alex Markov
Scott Duncan
Michael Balchanos
Woong Je Sung
Dimitri N. Mavris

Aerospace Systems Design Laboratory
School of Aerospace Engineering
Georgia Institute of Technology
Atlanta, GA 30332, USA

Margaret L. Loper
Elizabeth Whitaker
Michael Riley

Information and Communications Laboratory
Georgia Tech Research Institute
Georgia Institute of Technology
Atlanta, GA 30308, USA

ABSTRACT

The concept of Smart Cities is introduced to deal with the problems and challenges that come along with the growth of urbanization. Among all the challenges cities face, one fundamental concern is the safety of its citizens, especially during large sporting events. To test how smart capabilities can affect the evacuation process, an agent-based model was developed of the Georgia Tech football stadium. The model was developed to evaluate possible social behaviors during an evacuation, as well as produce the metrics of interest. To better visualize the results, a user interface was developed that includes a scenario evaluator and a decision making tool.

1 INTRODUCTION

The urban environment is becoming increasingly more connected and complex. In the coming decades, we will be surrounded by billions of sensors, devices, and machines, commonly called the Internet of Things (IoT). Cities and urban areas that benefit from the IoT are commonly referred to as Smart Cities. Smart Cities refer to urban areas that use different types of data to manage assets and resources efficiently by bringing together technology, government and society (Bawany and Shamsi 2015). By making cities "smart", the issues associated with urbanization growth can be addressed in a more efficient way, and solutions can be implemented more quickly (International Telecommunication Union 2012). Smart cities allow for improvements in many areas, including surveillance of threats, situational awareness, and security, which are exactly what's needed to avoid severe accidents that may happen in cities. By bringing these two concepts together, the internet of things for smart cities can help create safer, more efficient cities by transforming infrastructures, buildings and services with loT solutions (Vermesan and Friess 2013). The goal of this research is to develop a layered model of an urban emergency event, and determine the influence of smart devices and human first responders on the ability to evacuate a stadium.

2 METHODOLOGY

2.1 Problem Scoping and Framework

Atlanta is hosting 2019 Super Bowl at the Mercedes-Benz stadium, so we decided to model an emergency event at a football game. The Mercedes-Benz stadium is relatively new, which means access to data is challenging. Therefore, we decided to model the Georgia Tech (GT) Bobby Dodd stadium instead. The project objectives were to develop an agent-based model of a specific emergency event at the stadium, generate a variety of cases to test smart capabilities that assist evacuation of the stadium, and create a user interface to visualize the results. Figure 1 shows the framework. For the stadium, historical data is imported to calibrate the model. Smart capabilities are implemented to improve communication by identifying the emergent behaviors and the enablers. The visualization tool allows data to be displayed clearly, and can be used to compare multiple model runs. A user input interface is included in the visualization tool to set initial conditions and increase user-interactivity. The stadium model and the visualization tool are integrated to specify metrics of interest. The layout of the visualization tool also includes a decision-making tool, which provides users with a rank order of best cases with respect to desired metrics and presents the user with alternatives so they can make better decisions.

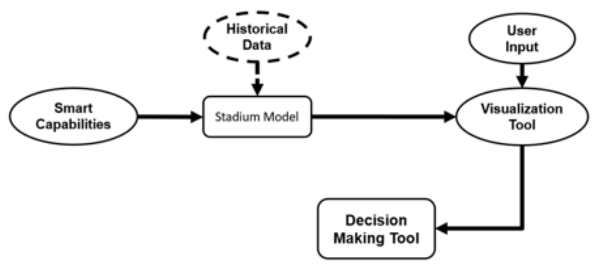

Figure 1: Framework of model.

2.2 Need for Agent Based Model and Interface Tool

The two main parts of the model are the environment, which is the Bobby-Dodd stadium, and the objects, which include the people attending the game and a suspicious package. We decided to use an agent-based model (ABM) to develop our simulation. ABMs are a class of computational models for simulating the actions and interactions of autonomous agents with the ability to assess their effects on the system as a whole. Advantages such as using a natural description of an event, giving insights into causes of phenomena and providing a framework for testing strategies are what we need to solve our problem more efficiently.

We identified several candidates ABM frameworks for our project, including Anylogic, AgentSheets, Repast, Cougaar, OpenStarLogo and NetLogo. Since we only had two semesters to complete the project, we did not have enough time to learn a new and complex computer language for coding and implementing the agent-based model. Instead we needed a tool that was relatively easy to learn. We also wanted a tool that would have a strong capability of modeling social behavior of people. Based on these requirements, we chose NetLogo to be our tool for implementing agent-based model.

The next step was to conceptually build our scenarios for analysis. As we mentioned previously, we decided to focus on the safety issue and evacuation plans of a football stadium. We established a scenario as follows: a typical game day is occurring at Bobby Dodd stadium and there are a lot of attendees inside the stadium, then a suspicious package is detected and the audience is asked to evacuate the stadium. In our scenario we focused on the evacuation plan and route instead of casualty. The suspicious package did not explode, it was just a danger which people tried to avoid. Several important issues that we wanted to concentrate on while modeling included the location of the package, the way attendees received the message for evacuation, their knowledge of evacuation routes, and the issue that some attendees do not follow the correct evacuation instructions.

A high-quality visualization was necessary for our project because it would allow us to have much better understanding of the data and results, and give the user (i.e., the decision maker) the ability to quickly and easily access the results, do comparisons between cases, and make the best decision based on the results. Since NetLogo is not able to compare multiple sets of data, and it has very limited number of visualization types. The results from NetLogo simulation needed to be further post-processed, compared and analyzed, so we developed a separate user interface. Since evacuation rate was one of our major concerns, we needed the interface to display data over time, and we also needed to analyze and compare different datasets.

We needed the visualization tool to have the ability to include an input section that is linked to NetLogo parameters. In order to make our case more realistic, we also wanted a real map of the GT stadium to represent the background for our scenarios. After a literature review, we found that a single computer language would not accomplish all of our requirements. As a result, we combined different languages to meet our goal: Hypertext Markup Language, JavaScript, and Cascading Style Sheets. We chose three free and open-source libraries of these languages, including Bootstrap, Leaflet and MapQuest API.

2.3 NetLogo Model Details

In order to fully understand the NetLogo model, some details about ABM need to be addressed. There are three key features to ABMs: the environment, the agents, and the rules. The environment is where all of the agents will interact. This can be hard coded in or imported from an image, but the latter has issues with resolution. For the stadium evacuation, the GT Bobby Dodd stadium was the environment. The agents interact in this environment and can receive information from it as well as affect it. Agents can have many different types, or breeds, which allows the creation of subsets of agents. These subsets can have different properties or behaviors and follow different sets of rules. In order for the agents to perform actions and for the simulation to occur, rules need to be determined for the agents to follow. These rules give bounds to what agents do and can be set to mimic certain behavior.

To get a better understanding of what the agents will actually be doing in the simulation for this analysis, a closer look at the coding approach and logic is beneficial. The first step is to assign rules to the agents and in this case the first action the agents will perform is to observe their environment. This allows the agents to gather information which will be used to decide what they will do next. Once the agents have gathered some information, they use it to change their state, which is a set of variables and values inherent to all agents. Finally, the agents will perform an action based on their new state and this process will be iterated until all of the agents leave the stadium. Figure 2 below shows a visual representation of the coding approach used for the agent based model. There are some steps required before starting the simulation, during, and after, so a breakdown of these has been created and is shown in Figure 3 below.

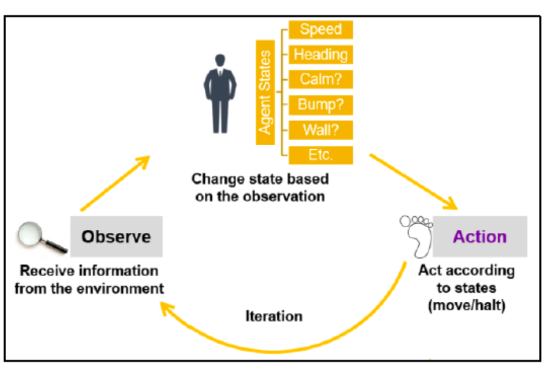

Figure 2: Coding approach for the agent based model.

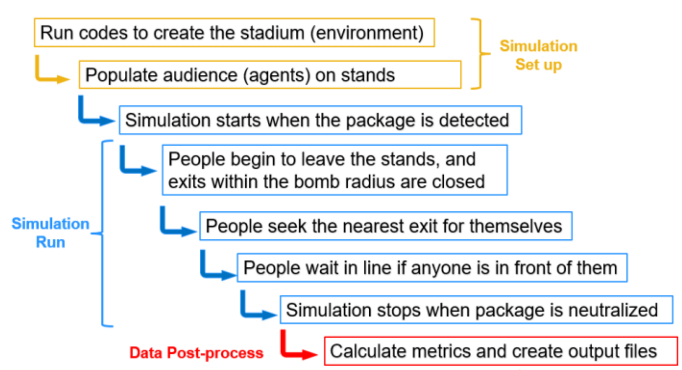

Figure 3: Steps of the NetLogo simulation process.

The simulation setup is always initiated before the simulation can begin. This set up generates the environment and populates the areas designated as stadium seats with agents. This also generates the other agents that are not necessarily attendees. A suspicious package is also generated before the simulation starts, although it is not required for the simulation to run properly. After the setup, the simulation will be initiated by the detection of the package and the agents will begin to leave the stands and evacuate. Any exits within the dangerous radius of the package will be closed and agents will not be able to evacuate through them. The agents will seek the nearest open exit, however there is a percentage of the agents which will exhibit chaotic or random behavior to simulate panic, and these agents will not seek the nearest exit. Agents also cannot move past other agents directly, but must wait until there is space in front of them before moving. This prevents multiple agents from standing in the same place or walking through other agents. Once the simulation plays out and all of the agents have evacuated, the package is assumed to be neutralized and the simulation is over. The final step is to post process the information and output the metrics of interest. In order to better visualize the simulation, an image of the stadium environment is shown below in Figure 4.

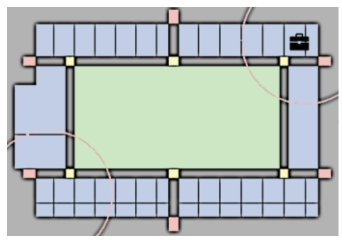

Figure 4: Layout of the stadium environment.

Now that model has been detailed and the simulation has been explained, a method for the analysis is required. This method should be a form of scenario analysis which looks at the impacts of smart technologies. The goal is to have the capability of comparing scenarios to a present day baseline, as well as identifying impacts and interactions of different technologies and parameters. Another important area to consider is the fact that some smart technology may be used in a negative or harmful way, and this should be addressed in these scenarios. To perform the analysis, three parameters were chosen to represent smart technologies: communication, misinformation (harmful technology), and first responders. The descriptions and real world equivalents of these parameters are given in Table 1 below.

Table 1: Descriptions of the parameters used in the scenario analysis.

	Communication	**Misinformation**	**First Responder**
Representation in Model	• Modeled as delay in evacuation • High communication corresponds to less delay • Three settings	• Incorrect location of package • Modeled as time step at which real location is identified • Three settings	• Purpose is crowd control • First responders are not physically modeled, only their effect of calming down the crowd • Three settings
Representation in Real World	• Verbal (face to face) • Broadcasting/Loud speaker • Digital message	• Hacking (spread of false info) • Cyber security system	• Real first responders • Emergency plan for evacuation

An important detail to note is that the first responders are not modeled as agents, but only their effect is modeled. The first responders act to control the crowds within the stadium and thus reduce the percentage of people which are in a panicked state. They are not present in the simulation but as more first responders are added, the percentage of panicked people at the initiation of the simulation is decreased, and the effects of this will be discussed later in the results and future work sections.

Each of the three parameters has three settings: low, medium, and high. It was decided to use these broad categories in order to parametrically explore the sensitivity of the model. Using the three parameters and three settings for each, a full factorial design space was created. The settings for the first responders are straight forward as they determine how many responders are present (0, 5, or 10). The settings for the communication and misinformation, however, have values of 0, 300, and 600 which correspond to ticks in NetLogo. Ticks are units of time - one tick is approximately equal to 1.6 seconds in the real world. If the delay is set to 300 ticks, this corresponds to a delay of eight minutes in real time. This allows for the model to be compared to real world events and also allows for calibration to test the evacuation time. A list of the parameter settings is shown in Table 2.

Table 2: List of parameters settings in NetLogo.

	Low	**Med**	**High**
Communication	**600**	**300**	**0**
First Responder	**0**	**5**	**10**
Misinformation	**0**	**300**	**600**

Note that this version of the simulation does not accurately capture the impact of communication, misinformation and number of first responders in an evacuation setting. Rather, it gives us the ability to explore model sensitivity, and focus on the design of the scenario evaluator and decision making tool functions of the user interface.

2.4 Interface Details

2.4.1 Metrics of Interest

We defined some metrics of interests that are key information that we will use to compare our different scenarios. We focused on the following parameters: evacuation time, number of people evacuated at each exit, and people in the danger zone. The evacuation time is represented by the number of people remaining in the stadium as a function of time. It is an important metric to evaluate an evacuation strategy. We will take into account different evacuation times: time to evacuate 100% of the stadium, 90% of the stadium and 80% of the stadium. For the number of people evacuated at each exit, we looked at the total number of people evacuated at each exit by the end of evacuation process. It is a useful metric for further analysis on traffic and congestion. Finally, we looked at the number of people in the danger zone. That metric is represented by the number of people within a certain radius of the suspicious package at time step t. It provides an estimate of the number of people that can get injured if the suspicious package explodes.

2.4.2 User interface – Scenario Evaluator

The User Interface has two main functions: Scenario Evaluator and Decision Making Tool. The goal of the scenario evaluator function, as shown in Figure 5, is to compare key metrics of a specific strategy to the baseline strategy. The baseline scenario is: Low Communication, Low Misinformation, Low number of First Responders. This function of the tool is composed by an *input* section that allows the user to choose initial conditions, an *output* section that allows data to display and information to be represented, and a *map* section which consists of a visual map of location or city which can be used to overlay data.

2.4.3 User interface – Decision Making Tool

The second function of the User Interface is to compare the entire set of scenarios against the metrics of interests. It consists in a general overview of all the scenarios that are ranked by parameter of interest, that can be chosen by the user. We defined three areas of focus for the rankings: focus on the evacuation time, focus on the number of people injured and focus on the cost of the rescue operation. These different foci are defined by changing the weight of the metrics of interests in the score calculation. For example, when the focus is done on the evacuation time, a higher coefficient is used when computing the score of the scenario. The scenario that obtains the minimum is ranked best strategy for the specific focus.

$$score = k_1 * (EvacTime100\% + EvacTime90\% + EvacTime80\%) + k_2 * NbPeopleInjured + k_3 * NbFirstResp$$

$$where\ k_1, k_2, k_3 = 1\ if\ focus\ on\ the\ parameter;\ \frac{1}{2} else$$

Parameters in the score calculation are scaled to work with numbers of the same magnitude. When attempting to assign a single "score" to solutions of a multi-objective problem, some weighting of preference is necessary across objectives. This is done simply in this paper, and the results are discussed, but there are other means of determining preferred solutions. The scoring method here is a placeholder until decision makers are engaged to determine a more tailored multi-objective function.

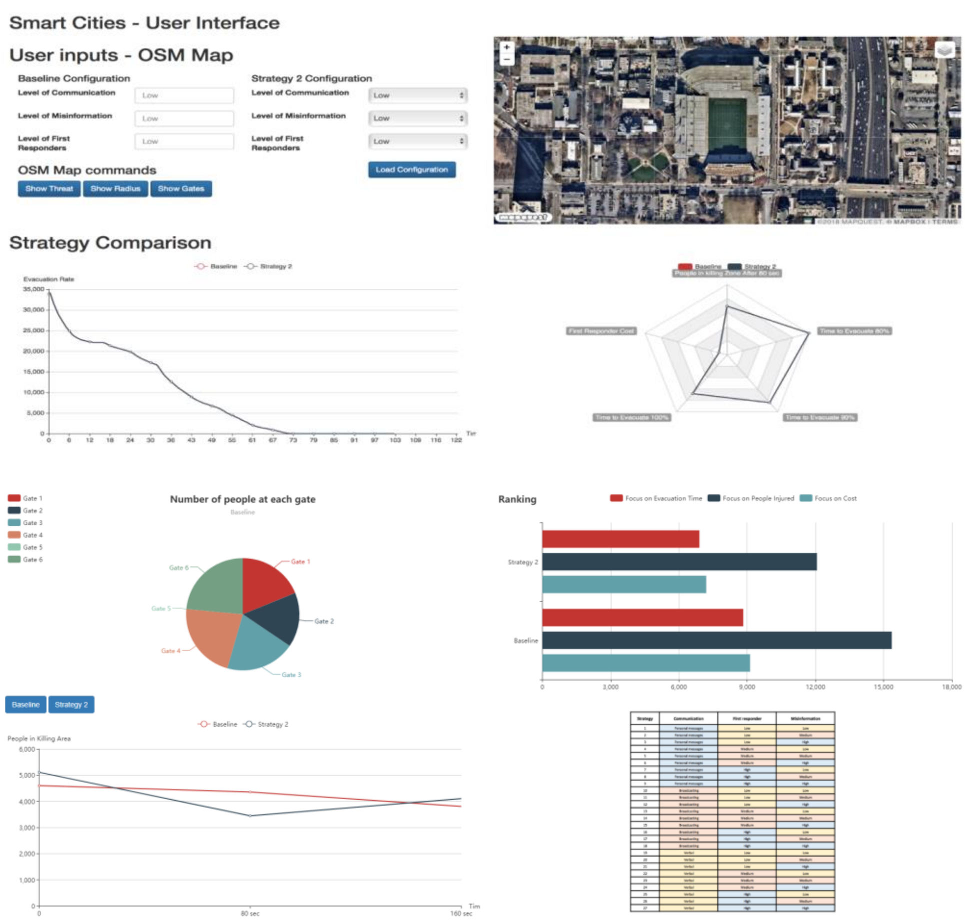

Figure 5: User interface scenario evaluator.

3 ANALYSIS AND RESULTS

3.1 Communication Impacts

In order to analyze the impacts of communication, three cases with different levels of communication, but with a constant level of misinformation and first responders were run in NetLogo. The three different communication levels were low, medium and high, which indicated 600, 300 and 0 ticks of delay in the simulation.[1] The total number of attendees in the stadium was 34,778. For all experiments presented in the paper, the initial number of people in the danger area is roughly ~4500 at t=0. The simulation results are shown as Figure 6, Figure 7, Table 3 and Table 4.

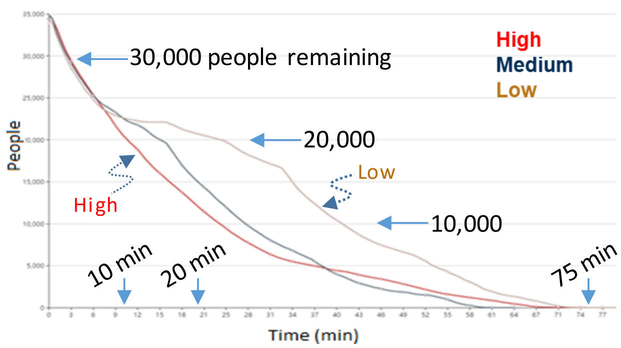

Figure 6: Number of attendees in stadium over time (H, M, & L levels of communication).

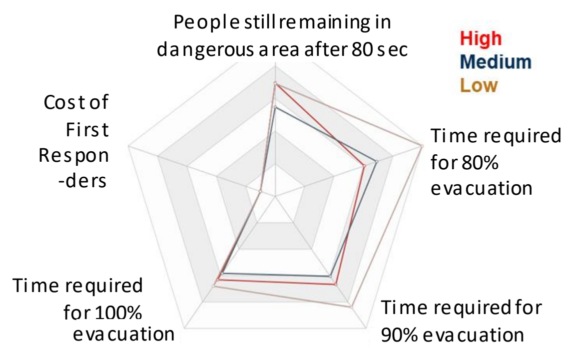

Figure 7: Evacuation metrics spider plot (H, M, & L levels of communication).

Table 3: Time to clear the danger area in stadium (communication).

	Low	Medium	High
80%	48	33	30
90%	57	41	46
100%	59	50	54

Table 4: People in dangerous area after 80 seconds (communication).

Low	Medium	High
4194	3445	4371

From the results we can observe that the time for evacuating all attendees in low communication case is the longest among all cases, which makes sense to us because low communication level means people are unable to get information very efficiently so there will be lags for operations, which will delay the evacuation time. And from Figure 6, we can see as the communication level becomes higher, the number of people in stadium decreases dramatically after 5 minutes. Consequently we come to the conclusion that better communication is required for more efficient evacuation.

However, at the same time we can observe that the case with highest communication level does not yield the shortest evacuation time, although it is still shorter than the lowest communication case. On the contrary, the case with medium level of communication has the best result in terms of evacuation time. If

[1] The use of ticks is to convey relative time impact (i.e., "which takes longer"), rather than focus on the total time. The conversion of 8 minutes to 300 ticks gives a relative magnitude.

we take a further look at the table recording the people in the danger zone after 80 seconds, we will find that the highest communication level has the most number of people in the danger, which is surely unexpected.

The main reason for the highest communication not having the best results could be that our agent-based model is still low-fidelity, so it misses some important conditions that are necessary for the model, which leads to the discrepancies between the intuition and the simulation results. Trying to include more features from reality and set rules based on them will be one of the major improvements for the model that can be done in future work.

In Table 4, the number of people in the danger area is not zero after 80 seconds; there will still be people inside after this time stamp. The 80 seconds is the time at which a measurement is taken, rather than the time required to evacuate the area. Additionally if a different suspicious package location is assumed, this will affect the size of the area and the time to evacuate it.

3.2 Misinformation Impacts

In order to analyze the impacts of misinformation, three cases with different levels of misinformation, but with constant levels of communication and first responders were run in NetLogo. The three different misinformation levels were low, medium and high, which indicated 0, 300 and 600 ticks for the fake package to be confirmed and neutralized in the simulation. The total number of attendees in the stadium was 34,778. The simulation results are shown as Figure 8, Figure 9, Table 5 and Table 6.

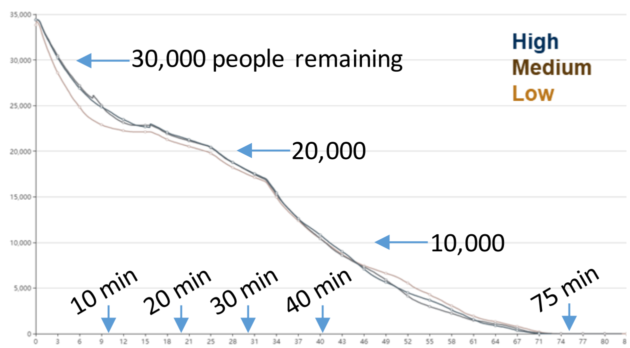

Figure 8: Number of attendees in stadium over time (H, M, & L levels of misinformation).

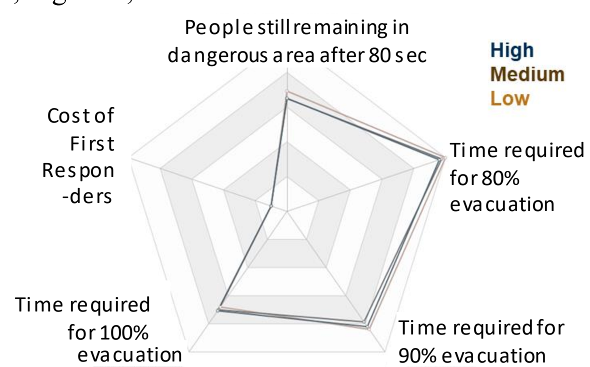

Figure 9: Evacuation metrics spider plot (H, M, & L levels of misinformation).

Table 5: Time to clear the danger area in stadium (misinformation).

	Low	Medium	High
80%	48	47	46
90%	57	61	56
100%	59	68	61

Table 6: People in dangerous area after 80 seconds (misinformation).

Low	Medium	High
4194	4082	4091

From Table 5 we can see the effects of misinformation. As the table shows, the lowest misinformation level yields the shortest evacuation time for the 90% and 100% cases. This matches our expectation, since misinformation sends people in wrong directions thus increasing evacuation time.

From the results we can observe that the misinformation level has slightly changed the evacuation time and the number of people inside dangerous area. By observing Figure 8, we can see three curves are actually pretty close to each other, and from Figure 9 we see that most metrics of interest do not change much due to different misinformation levels, both indicating that in our model, misinformation is not playing a major role. This is not satisfying, because we were expecting its effects to be more significant.

If we take a closer look at Table 5 we will find that the highest misinformation is not giving the longest evacuation time. And Table 6 is showing us that higher levels of misinformation will actually decrease the number of people in dangerous area. They both contrast with our expectation because we considered that misinformation, by misguiding people, would put more people in danger and delay the evacuation.

The cause for the unexpected results could be as same as the cause in section 3.1, which is that we are using a low-fidelity model. In our code logic, the only consequence by having misinformation is to temporarily close an exit (or several exits). This is apparently not enough for considering the consequence of misinformation in the real world. So coming up with more rules and conditions for misinformation in the future will improve the fidelity of our model, and produce better results.

3.3 First Responder Impacts

Finally, to analyze the impact of the first responders the setting was varied while keeping the other two parameters set to the “low” setting. The effect of the first responders is to reduce the percent of the population which is panicked. Less first responders results in a larger percent of panicked people. Although the numbers for this variable are 0, 5 and 10, they only represent 50%, 25% and 0% of the audience being panicked in the simulation run (in general, one first responder reduces 5% panic people). The results of the evacuation and the metrics for these cases are shown in Figure 10, Figure 11, Table 7, and Table 8 below.

These figures and tables show that although there is some change across the settings, the first responders do not seem to have a significant effect on any of the metrics other than the number of people in the dangerous area after 80 seconds. The impact also happens to increase this number which is the opposite effect of what was expected. This impact is also shown in the evacuation statistics because it takes a longer amount of time to evacuate with increased number of first responders. This effect is likely due to the responders not being modeled as agents themselves, but rather just their expected effect. Because the attendees all leave simultaneously with the maximum amount of responders (10), the exit gates become congested and this actually slows down the overall evacuation. The other settings did not change the evacuation time by more than one minute. As mentioned for the previous sections, more fidelity will be needed to truly understand why the results have come out this way and how to more accurately model human behavior, but this is discussed later for future work.

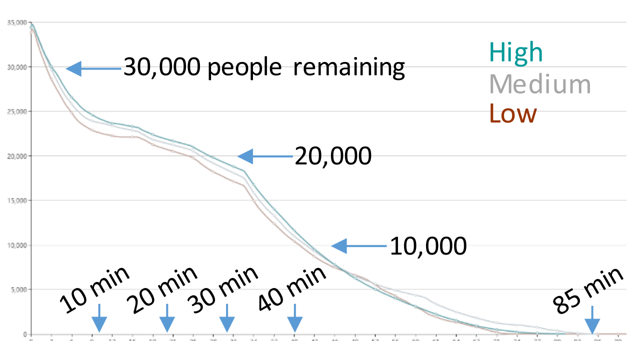

Figure 10: Number of attendees in stadium over time (H, M, & L levels of 1st responders).

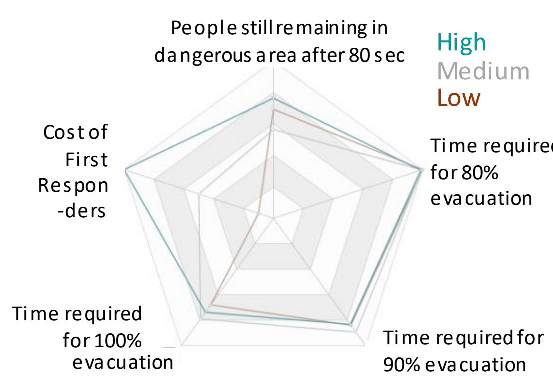

Figure 11: Evacuation metrics spider plot (H, M, & L levels of 1st responders).

Table 7: Time to clear the danger area in stadium (1st responders).

	Baseline	Medium	High
80%	48	48	48
90%	57	57	57
100%	59	59	63

Table 8: People in dangerous area after 80 seconds (1st responders).

Baseline	Medium	High
4194	4352	4796

3.4 Decision Making Tool

Assessing the results with the Decision Making tool, we obtained that the best scenario was a Medium level communication, No misinformation and Less first responders, which is Strategy 10 shown in Figures 12, 13 and 14. This strategy obtained the minimum score for the three ranking scores that are defined in section 3.3. These are not the results that we were expecting, but this is primary due to the fact that our model still need to be improved and validated against real data.

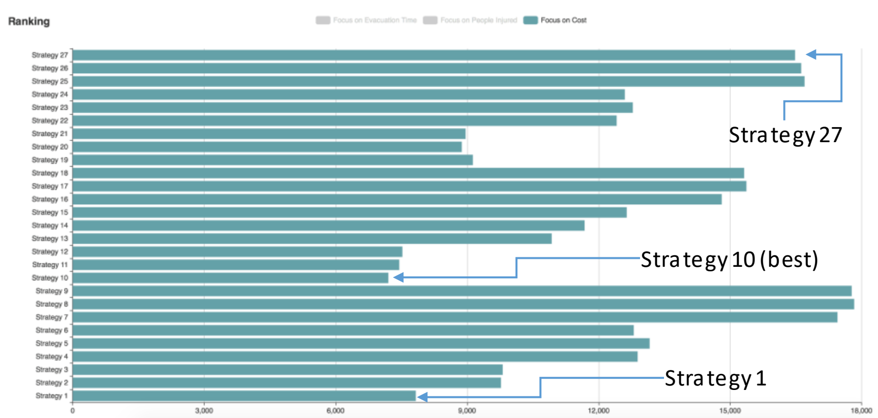

Figure 12 : Scenario ranking: focus on operation cost.

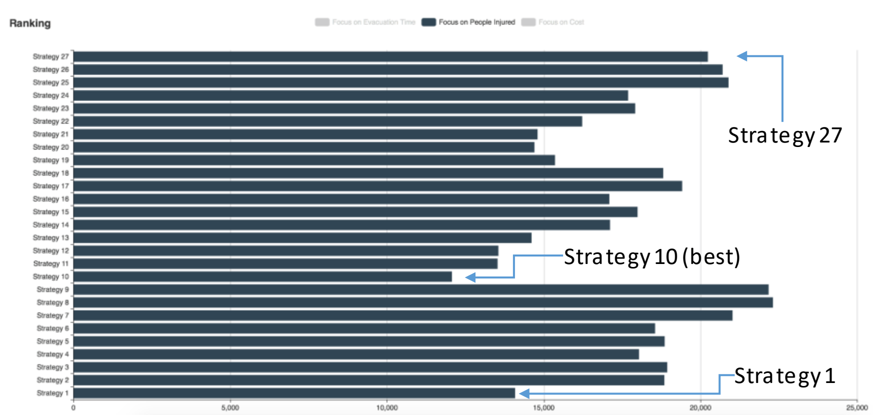

Figure 13: Scenario ranking: focus on people injured.

Figure 14 : Scenario ranking: focus on evacuation time.

4 CONCLUSIONS AND FUTURE WORK

For this project, a tool has been created for the purpose of assessing the impacts of technology on emergency evacuation scenarios. This tool is capable of comparing scenarios to a present a baseline as well as ranking the scenarios based on user preference. The scenarios have been defined to integrate smart capabilities, those can be useful but also harmful. The tool uses data from an agent based model which was developed to model the GT Bobby Dodd stadium. The end goal is to expand this agent based model urban environments, with an intermediate objective to scale up our current model for the Mercedes-Benz Stadium where will the 2019 Super Bowl will be held. Increased fidelity and more accurate human behavior will be needed in the future, to extend the current model and make it more realistic.

Some emergent behavior such as people getting stuck in queues and an increase in evacuation time for increased communication will have to be corrected because they do not model real behavior accurately enough. A different evacuation strategy than the delay may also be required for this because looking at Figure 6, Figure 8, and Figure 10, there is a point where the evacuation rate decreases due to congestion. This result is due to the agents in the danger zone leaving quickly, but eventually the exits become clogged, slowing the evacuation. Once the exits begin to free up, people evacuate faster. Another area for improvement is communication between agents, because this will impact the information that the agents receive, changing the way they escape. The agents may also interact with the environment in a new way. For example, if a sensor, represented by an agent, were placed at an exit, that exit could communicate with attendees if they should exit or find another route. This is also an effective way to compare the baseline scenario to a smart stadium or smart exit scenario which has some ability to send messages. Finally, as mentioned before, more agent breeds (types) and states will need to be introduced to capture a larger variety of behavior and produce more accurate simulations.

The results of the model do not necessarily agree with what was expected based on the inputs. This is likely due to lower fidelity of the model. In order to decrease the uncertainty, more detail will have to be added into the agent based model. This can, and should, be done to several areas of the model, the first of which is the diversity of agents. Some key features the need to be added for these agents will be a variety of speeds, varying levels of panic or calmness, and new states to name a few. Possibly the most important type of agent that needs to be introduced is one for the first responders. Because only their effect was modeled, not the responders themselves, the model was not significantly impacted and thus their effect was minimal. To extend on that idea, agents will also need to interact more with each other to try and capture more accurate human behavior. This can be done through new states and updated logic when looking for an exit. Along with new states, new behaviors will need to be introduced to the agents. Currently, the behaviors are limited and this is likely causing a less accurate representation of the movement of the attendees. Lastly, validation is part of future work. The scope of this work was to test out the behavior of the model and discuss the implications to support decisions, including how the user interface serves as a link from simulations to the exploration of scenarios.

REFERENCES

Bawany, N.Z. and J.A. Shamsi. 2015. Smart city architecture: Vision and challenges. International Journal of Advanced Computer Science and Applications, 6(11).

Becket, S. 2017. (October 4, 2017). "More details revealed about Las Vegas shooter's arsenal of weapons". CBS News. Retrieved from https://www.cbsnews.com/news/more-details-revealed-about-las-vegas-shooters-arsenal-of-weapons/.

Department of Homeland Security (DHS). 2015. "Securing the Super Bowl", retrieved from https://www.dhs.gov/blog/2015/01/30/securing-super-bowl.

International Telecommunication Union. 2012. "Overview of the Internet of things".

McLaren, D. and J. Agyeman. 2015. Sharing cities: a case for truly smart and sustainable cities. MIT Press.

Nations, U. 2014. World urbanization prospects: The 2014 revision, highlights. department of economic and social affairs. *Population Division, United Nations*.

Vermesan, O. and P. Friess, eds., 2013. Internet of things: converging technologies for smart environments and integrated ecosystems. Aalborg, Denmark: River Publishers. ISBN 978-87-92982-96-4.

ACKNOWLEDGMENTS

The authors would like to thank the Georgia Tech Research Institute (GTRI) for supporting this research under the Distributed Urban Sensing Technologies (DUST) strategic initiative.

AUTHOR BIOGRAPHIES

WENXIN ZHANG, ALEX MARKOV, VINCENT TERRIER, and XIAOSHU FEI are all Graduate Research Assistants in the Aerospace Systems Design Laboratory (ASDL) at Georgia Tech (GT). Together, they developed the agent-based simulation of the stadium evacuation. Their email addresses: stanleyzhang@gatech.edu, markov@gatech.edu, vincent.terrier@gatech.edu, xiaoshufei@gatech.edu.

SCOTT DUNCAN, MICHAEL BALCHANOS, and WOONGJE SUNG are all Research Engineers in the ASDL and lead developers under GT's Smart Campus Initiative, where visual analytics, data science, and modeling and simulation have been applied to improve campus assessment, operations, and planning. Email: duncan@gatech.edu; michael.balchanos@asdl.gatech.edu; woongje.sung@asdl.gatech.edu.

DIMITRI MAVRIS is the S.P. Langley Distinguished Regents Professor, Director of the ASDL, and Principle Investigator for the GT Smart Campus Initiative. Email: dimitri.mavris@ae.gatech.edu.

MARGARET L. LOPER is a Principal Research Scientist at the Georgia Tech Research Institute (GTRI), where she leads research on simulating smart cities. Email: margaret.loper@gtri.gatech.edu.

ELIZABETH WHITAKER is a Principal Research Engineer at GTRI, where she leads research in cognitive systems. Email: elizabeth.whitaker@gtri.gatech.edu.

MICHAEL RILEY is a Research Scientist at GTRI, where he develops software for agent based models. Email: michael.riley@gtri.gatech.edu.