USING REINFORCEMENT LEARNING FOR ACTIVE SHOOTER MITIGATION

by

Robert E. Bott

A Dissertation

Submitted to the Faculty of Purdue University

In Partial Fulfillment of the Requirements for the degree of

Doctor of Philosophy

Department of Technology

West Lafayette, Indiana

December 2021

THE PURDUE UNIVERSITY GRADUATE SCHOOL
STATEMENT OF COMMITTEE APPROVAL

Dr. J. Eric Dietz, Chair

Department of Computer and Information Technology

Dr. Eric T. Matson

Department of Computer and Information Technology

Dr. Baijian Yang

Department of Computer and Information Technology

Dr. Charles E. Anklam III

United States Marine Corps

Dr. Braiden M. Frantz

United States Marine Corps

Approved by:

Dr. Kathryne A. Newton

In Every Clime and Place

(Even in comfort, riding a desk with coffee in hand)

ACKNOWLEDGMENTS

Nothing is ever achieved on your own. From the smallest to the largest accomplishment, it is the efforts of others that guide, mentor, and drive us to our goals and beyond. With that in mind, I would like to acknowledge everyone in my life who has influenced me both positively and negatively. You are the reason I am here; You are what made me get up every day, no matter the hour, to either fight you, in spirit or until death, or to serve you and make content. My efforts might be my own, but the motivation and drive were always you. Always you.

I would like to mention several people in my recent life that stand out in providing motivation and much-needed assistance in skillsets and knowledge I lacked. My committee Chairman, Dr. J. Eric Dietz, has always been there to advise and assist, and for that, I am grateful. He has let me do what I thought was right without allowing me to chase the rabbits too deeply down those holes. Thank you, Sir! Other members of my committee deserve more thanks than I can give for providing guidance on their specific expertise and other areas of research and dissertation writing. Dr. Anklam III., Dr. Frantz, Dr. Matson, and Dr. Yang, thank you, gentlemen, for your support. My work is better because of you.

Many thanks and my appreciation go out to my fellow students who have taught me much. Braiden, thank you for being a calm and measured ally throughout our time together, working as two Marines should: As a team, with a mission in mind and an end state to reach. Krassi, thank you for your instruction on AnyLogic, and much comic relief when the times called for it. Jae, the “maker of models” who set us all on a reliable path of quality research at the Purdue Homeland Security Institute by creating a genuinely stunning model we could all base our work on. To all the others at the PHSI who have crossed paths with me, thank you. I salute you all!

I could not have done what I did so quickly and efficiently without the Pathmind system and the people at that company who assisted. Ty, for letting me use the system early on deserves thanks. Sahar for helping me to better the reinforcement learning side of things. Above all, Edward who was there for me to bother without fail and who greatly assisted in tracking down bugs in my model. Thank you all.

My thanks and appreciation go to Sarah, the princeps editor, for helping me make this document, and many others, appear professional and readable. I am confident that you have saved many readers from severe headaches.

Finally, thank you Jessi for being a steady partner in all things, and for convincing me of the value of the Oxford comma (somewhat). I guess The Brits do have something to say about the English language that is worth considering.

TABLE OF CONTENTS

LIST OF TABLES 9
LIST OF FIGURES 10
GLOSSARY 12
LIST OF ABBREVIATIONS 14
ABSTRACT 15
CHAPTER 1. INTRODUCTION 16
1.1 The Active Shooter Problem 16
1.1.1 Open-Air Venues and Spaces 18
1.1.2 Machine Learning 19
1.2 Statement of the Problem 19
1.3 Significance 20
1.4 Research Questions 21
1.5 Purpose 21
1.6 Scope 22
1.7 Assumptions 22
1.8 Limitations 23
1.9 Delimitations 23
CHAPTER 2. REVIEW OF LITERATURE 25
2.1 Methodologies 25
2.2 Active Shooting Incident Terms and Definitions 27
2.3 Active Shooting Incident Statistics 30
2.4 Active Shooting Incidents and the Evolution of Responder Training and Tactics 33
2.4.1 Open-Air Venues and Open Spaces 42
2.5 Agent-Based Modeling and Simulation 46
2.5.1 Active Shooting Incident Research using Agent-Based Modeling and Simulation .. 49
2.6 Machine Learning 52
2.6.1 Reinforcement Learning 54
2.6.2 Deep Reinforcement Learning and Agent-Based Modeling and Simulation 55
2.7 The State of ASI Research Using ABMS and DRL 58

2.8 Summary ... 59
CHAPTER 3. RESEARCH METHODOLOGY ... 60
3.1 Overview ... 60
3.2 AnyLogic Modeling Software ... 60
3.3 Pathmind Software ... 61
3.4 Gilroy Garlic Festival Model ... 63
3.4.1 AnyLogic Model Foundation ... 63
3.4.2 Foundational Model Design Details ... 64
3.4.3 Foundational Model Parameters ... 66
3.4.4 Foundational Model Validation ... 68
3.5 Expanded Model Design ... 70
3.5.1 Adjusted Model Parameters ... 70
3.5.2 Model Environment ... 71
3.5.3 Added Model Parameters ... 72
3.5.4 Dynamic Signs ... 73
3.5.5 Shooter Agent Design and Logic ... 76
3.5.6 Civilian Design and Logic ... 78
3.5.7 Experiment Screen and Setup ... 82
3.6 Pathmind Integration and Use ... 84
CHAPTER 4. EXPERIMENTION, ANALYSIS AND RESULTS ... 88
4.1 Overview ... 88
4.2 Baseline Experiments ... 88
4.2.1 Civilian Run to Closest Exit ... 88
4.2.2 All Signs Active ... 89
4.3 Pathmind Experimentation ... 90
4.3.1 Initial Experimentation ... 91
4.3.2 Final Experimentation to Select Policy ... 93
4.4 Monte Carlo Experiments with Chosen Policies ... 101
4.4.1 In Pursuit of Research Question One ... 102
4.4.2 In Pursuit of Questions Two and Three ... 106
4.4.3 Confirming Reduced Sign Scenarios ... 119

4.5 Research Question One ... 120
4.5.1 Answering Research Question One ... 120
4.6 Research Question Two ... 121
4.6.1 Answering Research Question Two ... 121
4.7 Research Question Three ... 122
4.7.1 Answering Research Question Three ... 122
4.8 Further Experimentation ... 123
CHAPTER 5. SUMMARY AND RECOMMENDATIONS ... 125
5.1 Overview ... 125
5.2 Significance of this Study ... 126
5.3 Future Technology Development ... 127
5.3.1 Dynamic Signage ... 127
5.3.2 Unmanned Aerial Vehicles ... 128
5.3.3 Sensors for Detecting and Locating a Shooter ... 129
5.3.4 Integrated Defensive System ... 130
5.4 Future Research ... 131
5.4.1 Improving the Current DRL Policy ... 132
5.4.2 Improving the Model ... 133
5.4.3 A Note on Fidelity and Abstraction ... 135
APPENDIX A. EXPERIMENT SUMMARY DATA ... 137
APPENDIX B. SPAWN RELATED SCENARIO RUN DATA ... 139
APPENDIX C. COMPLETE SOURCE DATA SETS AND MODEL ... 167
REFERENCES ... 168

LIST OF TABLES

Table 3.1 - Model Parameter Description 67

Table 3.2 - Model Validation Parameters 69

Table 3.3 - Adjusted Model Parameters 71

Table 3.4 - Added Model Parameters 73

Table 3.5 - Observations and Reward Metrics 86

Table 4.1 - Run to Closest Exit Results 89

Table 4.2 - All Signs Active Results 90

Table 4.3 - Pathmind Trained Policy Results 102

Table 4.4 - Results Comparison Summary 105

Table 4.5 - All Signs and Locations Pre-Reduction 107

Table 4.6 - North Spawn Case Removed Signs 112

Table 4.7 - South Spawn Case Removed Signs 113

Table 4.8 - East Spawn Case Removed Signs 114

Table 4.9 - West Spawn Case Removed Signs 115

Table 4.10 - Reduced Sign Casualty Comparison 120

Table 4.11 - Policy Comparison in Differing Environment 123

LIST OF FIGURES

Figure 2.1 - Total Active Shooter Incidents 2000-2019 32
Figure 2.2 - Mass Shootings 1966-2000 35
Figure 3.1 - Example Pathmind Experiment Setup 63
Figure 3.2 - Model with 17 Exits 65
Figure 3.3 - Adjusted Model Environment 72
Figure 3.4 - Base Dynamic Sign Placement 75
Figure 3.5 - Dynamic Signs During Simulation Run 76
Figure 3.6 - Shooter Logic 77
Figure 3.7 - Shooter Agent Schematic 78
Figure 3.8 - Civilian Logic 79
Figure 3.9 - Civilian Logic Workflow Diagram 80
Figure 3.10 - Civilian Agent Schematic 82
Figure 3.11 - Experiment Screen 83
Figure 3.12 - Experiment Screen Running 84
Figure 3.13 - Pathmind Reinforcement Learning Values 87
Figure 4.1 - Pathmind Experiments List 92
Figure 4.2 - Experiment One Parameters 93
Figure 4.3 - Experiment Two Parameters 93
Figure 4.4 - Experiment Three Parameters 93
Figure 4.5 - Experiment Four Parameters 94
Figure 4.6 - North Spawn Experiment One Metrics 94
Figure 4.7 - North Spawn Experiment Two Metrics 94
Figure 4.8 - North Spawn Experiment Three Metrics 95
Figure 4.9 - North Spawn Experiment Four Metrics 95
Figure 4.10 - South Spawn Experiment One Metrics 96
Figure 4.11 - South Spawn Experiment Two Metrics 96
Figure 4.12 - South Spawn Experiment Three Metrics 97

Figure 4.13 - South Spawn Experiment Four Metrics 97
Figure 4.14 - East Spawn Experiment One Metrics 98
Figure 4.15 - East Spawn Experiment Two Metrics 98
Figure 4.16 - East Spawn Experiment Three Metrics 98
Figure 4.17 - East Spawn Experiment Four Metrics 99
Figure 4.18 - West Spawn Experiment One Metrics 99
Figure 4.19 - West Spawn Experiment Two Metrics 99
Figure 4.20 - West Spawn Experiment Three Metrics 100
Figure 4.21 - West Spawn Experiment Four Metrics 100
Figure 4.22 - North Spawn Results Comparison 103
Figure 4.23 - South Spawn Results Comparison 103
Figure 4.24 - East Spawn Results Comparison 104
Figure 4.25 - West Spawn Results Comparison 104
Figure 4.26 - Base Sign Placement Before Individual Removal 116
Figure 4.27 - North Spawn Case Removed Signs 117
Figure 4.28 - South Spawn Case Removed Signs 117
Figure 4.29 - East Spawn Case Removed Signs 118
Figure 4.30 - West Spawn Case Removed Signs 118

GLOSSARY

For this research study, the below terms are used throughout this work.

- Active Shooter: As per the FBI (Federal Bureau of Investigation, 2021) "… one or more individuals actively engaged in killing or attempting to kill people in a populated area. Implicit in this definition is the use of one or more firearms." (p. 1).
- Active Shooting Incident: An incident where an active shooter perpetrates his crimes.
- Agent-Based Modeling and Simulation: Computerized model used to gain insight into a complex system's behavior, using individual agents within the simulation (Bandini et al., 2009).
- Discharge: A term used to imply that a firearm has been operated.
- Machine Learning: The science (and art) of programming computers so they can *learn from data* (Géron, 2019).
- Mass Casualty Shooting: The murder of three or more individuals (United States Department of Justice Office for Victims of Crimes, 2017).
- Mass Casualty Incident: DeNolf and Kahwaji (DeNolf & Kahwaji, 2020) define it as "an event that overwhelms the local healthcare system, where the number of casualties vastly exceeds the local resources and capabilities in a short period of time."
- Mass Killings: Three or more killings in a single incident (Krouse & Richardson, 2015)
- Mass Murder: Multiple homicide incidents in which four or more victims are murdered within one event and in one or more close locations (Krouse & Richardson, 2015)
- Mass Shooting: An event where multiple individuals fall victim to an active shooter.

- Open-Air Venue: An event where patrons spend the majority of their time in the open or outdoors and not within a fully enclosed structure.

- Reinforcement Learning: One of the major areas of machine learning, defined as learning the optimal behavior in a specific environment to obtain the maximum reward (*What Is Reinforcement Learning?*, 2021).

- RUN.HIDE.FIGHT®: Actions that can be taken by individuals to safeguard their lives during an active shooting incident (City of Houston, 2012).

LIST OF ABBREVIATIONS

ABM	Agent-Based Model
ABMS	Agent-Based Modeling and Simulation
ALERRT	Advanced Law Enforcement Rapid Response Training
ASI	Active Shooting Incident
CCW	Concealed Carry Weapon
DRL	Deep Reinforcement Learning
FBI	Federal Bureau of Investigation
FOV	Field of View
RHF	Run. Hide. Fight.
RL	Reinforcement Learning
SRO	School Resource Officer
SWAT	Special Weapons and Tactics
UAV	Unmanned Aerial Vehicle

ABSTRACT

This dissertation investigates the value of deep reinforcement learning (DRL) within an agent-based model (ABM) of a large open-air venue. The intent is to reduce civilian casualties in an active shooting incident (ASI). There has been a steady increase of ASIs in the United States of America for over 20 years, and some of the most casualty-producing events have been in open spaces and open-air venues. More research should be conducted within the field to help discover policies that can mitigate the threat of a shooter in extremis. This study uses the concept of dynamic signage, controlled by a DRL policy, to guide civilians away from the threat and toward a safe exit in the modeled environment. It was found that a well-trained DRL policy can significantly reduce civilian casualties as compared to baseline scenarios. Further, the DRL policy can assist decision makers in determining how many signs to use in an environment and where to place them. Finally, research using DRL in the ASI space can yield systems and policies that will help reduce the impact of active shooters during an incident.

CHAPTER 1. INTRODUCTION

1.1 The Active Shooter Problem

Active Shooting Incidents (ASIs) are an increasing problem within the United States of America. They show no signs of abating, but rather since the year 2000, are increasing in frequency and overall lethality (Federal Bureau of Investigation, 2021). The most notorious and impactful ASI was the Columbine High School shooting on April 20th, 1999. This event changed the nation's perception of "active shooting" and brought the horrors of this kind of massacre home to millions. The death of 13 innocents at the hands of fellow students was shocking to behold for a country that was, at that time, not used to seeing this brutality play out within its own borders. By the year 2019, 333 ASIs occurred within schools, businesses, and public places and organizations, as well as many other locations, according to the FBI (Federal Bureau of Investigation, 2021). This is a staggering amount of violence perpetrated upon innocent individuals in an otherwise safe and secure environment. Few expect to have to confront this kind of violence by simply going about their lives.

The Columbine High School shooting also dramatically changed police and first responder tactics. During Columbine and before, first responders would secure a perimeter around an incident site and wait for special weapons and tactics (SWAT) teams to arrive to either negotiate with or engage the shooter. Unfortunately, during the Columbine incident, these tactics only gave the shooters more time to conduct their barbaric actions and resulted in more casualties, as one study in particular has highlighted (Lee, 2019). Modern first responder tactics assume that every second the shooter is still active puts innocents in danger. Therefore, police are trained to immediately engage the perpetrator by any means available. According to the FBI, 69.8% of incidents end within five minutes or less, and approximately half of those end in under two minutes (Blair, J.

Pete & Schweit, Katherine W., 2014). Time is of the essence, and this realization has had an enormous impact on how law enforcement responds to active shooting incidents. It has also opened discussions and research studies into different response methods, such as armed civilians with concealed carry weapons (Bott et al., n.d.). Though these tactics often result in fewer civilian casualties, they also put first responders in more danger (Bott et al., n.d.). Training methodologies and programs for civilians and law enforcement officers have been developed and tested in the intervening years to increase the chances of surviving an incident unharmed (City of Houston, 2012) (Mallonee, 2017).

Though ASIs have been burned into the conscience of the average citizen as occurring primarily in schools, many other locations are the target of killers. Private and government offices, churches, open-air venues, and even military installations have been the target of active shooters. The FBI has classified 12 different location types, of which out of 333 shootings between 2000-2019, only 62 (18.6%) occurred in K-12 or higher educational institutions (Federal Bureau of Investigation, 2021). This spread among a diverse group of locations indicates that authorities and planners cannot focus in on one or two problem areas. This further complicates the effective adoption of policies meant to protect the public and produce a more proficient first responder. It also speaks to the mental state of shooters that they do not simply show up at one or two particular target types. One can certainly claim that an active shooting can occur anywhere people live and go about their business. As such, policy makers must adopt flexible policies that can be used for a variety of situations and that can be adjusted to fit a particular venue better, with little effort on the part of organizers and responders. It also speaks to the need for proper training of both civilians and first responders to be able to protect themselves and others during an ASI. This research seeks to inform policy makers so they may develop better policies and technologies to guard against the active shooter threat.

1.1.1 Open-Air Venues and Spaces

Open-air venues and spaces offer a unique problem to first responders and civilians. The researcher seeks to focus on open-air venues and spaces since current research has not given much attention to these environments. By and large, open spaces and venues have been ignored by researchers in favor of enclosed spaces, such as office buildings and schools. This is despite the FBI stating that 50 out of 333 ASIs (15%) have occurred in what it classifies as open spaces, which is second only to businesses that are open to pedestrian traffic (96 instances). Other unspecified locations that might also contain open spaces and venues, are malls (10 instances) and military properties (nine instances) (Federal Bureau of Investigation, 2021). The FBI states that malls contain these open spaces though the sizes are not specified. Military installations, airfields, and port facilities are also mentioned as places ASIs occurred (Federal Bureau of Investigation, 2021). This illustrates the increase in occurrences of these shooting events in open spaces and venues, creating more impetus to conduct research specifically for those environments.

Even though current research for ASIs in open spaces and open-air venues is minimal compared to enclosed spaces, one particular author has contributed significantly to the field. The study looked at various open-air venues, such as the Las Vegas concert shooting and the Garlic Festival in Gilroy, CA (Frantz, 2021). Frantz (2021) built a computerized model of the Garlic Festival environment to produce data that supports better defensive policy in open-air venues. It is a significant contribution to the field, and this researcher has chosen to build upon that study to find more data to support better policy. Building upon this recent dissertation will yield significant knowledge on how to better defend against active shooters in open-air venues and open spaces in general.

1.1.2 Machine Learning

The author has chosen to focus his efforts not just on computerized modeling, but also on using machine learning (ML) to support his efforts. Specifically, the author will use deep reinforcement learning (DRL) to train various agents within the agent-based model. It is an important step towards using new and emerging technology to support ASI research that leads to the discovery of new and novel policies and technologies to defend against violence.

1.2 Statement of the Problem

This work addresses the problem of the vulnerability of open-air venues and open spaces to active shooters. Open spaces and venues like concerts, festivals, and various theme parks are often temporary and void of permanent and hard structures (Frantz, 2021). This forces the majority of civilians to congregate in open areas, often in groups, presenting lucrative targets to any active shooter. Further, without hard structures, these venues lack any protection for civilians against gunfire (Frantz, 2021). Even indoor venues with large open areas, such as a concert, present the same problems in the protection of patrons. Given these facts, it is clear how open-air venues and open spaces can be especially vulnerable to active shooters. One particularly lethal attack on an open-air venue was the Las Vegas Route 91 Festival shooting, where approximately 868 individuals sustained injuries during the shooting, and 58 of those were murdered (Joseph Lombardo, 2018). Though we will never know for sure, one can assume that the shooter chose the location based on the openness of the venue and the availability of his elevated ambush site in an adjacent building. This shooting also demonstrates the danger of patrons receiving injuries as a result of stampeding, crushing, or bludgeoning by panicked individuals. During the Las Vegas shooting, 456 victims were injured by other means than gunshot or shrapnel (Joseph Lombardo, 2018). Another issue concerning open-air venues is that many are temporary and "pop up" in various places throughout the country. Frantz (2021, p. 16) brings particular attention to this issue with his work on the Gilroy Garlic Festival shooting in 2019 and states:

> While some open-air venues like amusement parks are in the same location year-round, many are typically pop-up style, which are not permanent and only in place for a short period of time. This poses a unique problem set for event planners in regard to protection from a potential active shooter attack. Many open-air venues lack permanent structures that could shield victims from shooter gunfire and patrons will be unfamiliar with all available exits. (p. 16)

Without permanent structures and long-rehearsed security procedures, such events are more vulnerable to active shooters above other environments.

1.3 Significance

As previously mentioned, active shooting incidents (ASIs) are rising throughout the United States of America. The lethality to innocents is also increasing, while effective policies reducing lethality are limited. Specifically for open-air venues and other open spaces, there are few coherent and effective policies that decision makers can rely on at this time. Though there are actions that a civilian can take, such as running from the shooter, hiding from the shooter, or fighting the shooter, these often fail in extremis, particularly in open-air venues (Frantz, 2021). This is because of the many new and broader-ranged variables introduced in open spaces and venues, such as environment, logistics, and procedures.

Further, most venues restrict the ability for civilians to be armed for self-defense, despite research showing that having responders present immediately when the incident begins reduces overall casualties (Bott et al., n.d.). In general, most ASIs will end within five minutes (Blair, J.

Pete & Schweit, Katherine W., 2014). This leaves very little time for civilians or first responders to act appropriately while under duress. Research must be done to limit the thinking required for individuals to survive under stress, and to aid in taking appropriate actions during an ASI.

This research seeks to discover and test new methods, tools, and technologies that might help decision makers produce new policies or improve existing ones to help save lives.

1.4 Research Questions

The following research questions are posed in this work:

1. What impact does the application of reinforcement learning have on the number of casualties during an active shooting incident?
2. How many reinforcement learning controlled dynamic signs are needed to reduce casualties within an environment?
3. At what locations should reinforcement learning controlled dynamic signs be placed to reduce casualties within an environment?

1.5 Purpose

The purpose of this work is to examine the viability of reinforcement learning (RL), together with an agent-based model (ABM), in testing new and novel technologies and processes that could save lives during an active shooting incident (ASI). The researcher will rely on the efforts of a previous scholar who used the Garlic Festival ABM to discover low-budget safety protocols that can be implemented throughout open-air venues (Frantz, 2021). Choosing to build and expand a model based on a historical event has yielded useful insights and allowed for the current author to continue where the previous work left off. Further, given the validation efforts of the previous researcher, the current author is in a fortuitous position to continue where the previous work left off.

1.6 Scope

This research ultimately is intended to discover means to reduce casualties in open spaces and open-air venues. It also uses tools and technology that, as far as the author knows, have never before been used together in the field, which is significant in itself. A previously built model was modified and adjusted to fit into the new technology framework. This model covers the Gilroy Garlic Festival event and had been validated by (Frantz, 2021). Only one model will be used to conduct all work, given that the effort required to produce multiple models and integrate reinforcement learning agents would far exceed a realistic workload for research of this nature.

Data will be produced by training reinforcement learning agents within the Pathmind and AnyLogic systems. AnyLogic is the premier modeling tool used by academia, industry, and government to create agent-based models to produce data for the analysis of real-world scenarios (AnyLogic, n.d.). Pathmind is a state-of-the-art deep reinforcement learning (DRL) tool that integrates with AnyLogic to allow the modeler to use the power of reinforcement learning for agent-based models. By running the trained agents within the model using Monte Carlo runs, the researcher can collect the necessary data.

1.7 Assumptions

The following assumptions will be stated for this research project:

1. The active shooter agent will engage the closest civilian within its field of view.
2. The location of the shooter is always known by the system.

3. Civilians and the shooter will have their movement hindered by structures in the environment.
4. Civilians will exit the festival grounds via one of the exits that surround the environment.
5. Dynamic signage shall guide civilians towards the most advantageous exit, away from the shooter.

1.8 Limitations

Limitations for this research are listed below:

1. Research does not include multiple agent-based models.
2. A finite number of dynamic signs and exits are present in the model.
3. The only function of the dynamic signs is to be activated or de-activated.
4. Civilian agents will not fight or hide, only run to safety.

1.9 Delimitations

Delimitations for this research are listed below:

1. The model was built using AnyLogic and the deep reinforcement learning policy trained using Pathmind.
2. The scenario used is a modified Gilroy Garlic Festival agent-based model.
3. Deep reinforcement learning trained policies will only affect a single agent that controls all dynamic signs.
4. No further validation occurred on the Gilroy Garlic Festival agent-based model.

5. Civilian patrons will move towards the farthest activated dynamic sign in their field-of-view until they detect an exit, then will leave the simulation to safety through this exit.
6. No responders are active in the scenario to confront the shooter.

CHAPTER 2. REVIEW OF LITERATURE

2.1 Methodologies

Active shooting incidents (ASIs) have no clear patterns as to where they occur, who the shooters might be, and who the intended targets are. Frantz (2021, p. 30) states that "The actions of the shooter are unpredictable …" and that "Individuals who orchestrate these crimes are struggling with mental health issues and have often dealt with trauma in their personal lives.". Though many studies give us some idea who the primary perpetrators might be, where they often do their evil deeds, and what weapons they use, there is no clear and discernable pattern. It is difficult to find certainty as to who might become a shooter and how to prevent them. Conducting research in an attempt to derive policies to prevent shootings is a noble undertaking and well worth the investment. Yet, we cannot ignore policies that help increase survival during and immediately after an incident. Even incidents that result in relatively few casualties have an enormously damaging impact on our national conscience, not to mention the impact on the victims and their families. This work, and the research it relies on, is focused on protecting people during an ASI.

Over the last few decades, new procedures and policies have been developed focusing on the responder and the civilian under threat. Since the Columbine High School massacre, police have changed their tactics significantly. They no longer set a cordon and wait the shooters out, seeking a negotiated solution. Police across the country are now trained to immediately run into the structure and engage the shooter as soon as possible to reduce civilian casualties (Frantz, 2021). This, of course, puts responders at higher risk, and much attention has been given to this problem by law enforcement organizations and researchers at various institutions (Texas State University, 2021). Also, private and local training is offered to law enforcement by various former law enforcement officers and military veterans (BRIGGS CORE DYNAMICS, n.d.). Other organizations and law enforcement have also teamed up to provide basic threat assessments for the security of schools, further hardening these potential targets (National Rifle Association of America, n.d.). For civilians, specific policies have been developed focused on the actions that they might take if they find themselves in an ASI. For example, the RUN.HIDE.FIGHT® method has been developed by Ready Houston (City of Houston, 2012). Another methodology focused on civilian response to an ASI is Avoid | Deny | Defend ™, which is part of the ALERRT system created in part by Texas State University (Texas State University, n.d.).

One of the major problems in discovering new policies is the ability to assess new tactics, techniques, and procedures that might affect higher survival rates for both responders and civilians during an ASI. It would be impossible to conduct a real event with real perpetrators and weapons within actual structures. Other methods such as live training with simulated ammunition can help in training and assessment. However, that type of training and evaluation will never come close to the terror and complexity of a real event. Also, this training often lacks the participation of civilians and is logistically difficult to organize. Of course, we can proceed with after-action reviews and determine what participants could have done differently to increase survival. However, this is also limited due to individual bias and a lack of certainty in the actions of those who were participants.

Thanks to modern technology and software, decision makers and researchers can now experiment with and discover new tactics, techniques, and procedures that yield beneficial policy. Tools like agent-based modeling and simulation (ABMS) and supporting technology such as the AnyLogic software system (AnyLogic, n.d.) can help in this experimentation. These technologies allow the discovery of new, high-impact policies without the risks and logistical challenges mentioned before. They provide the means to achieve whatever level of fidelity or abstraction is required to produce noteworthy results. Using this methodology is a time-tested approach that generates quality data for analysis, and it is the main approach in this work.

2.2 Active Shooting Incident Terms and Definitions

As of the writing of this work, there are numerous terms, definitions, and properties for what an active shooting incident (ASI) is composed of. These have been published by various governmental and non-governmental organizations, yet they do not paint a complete picture of what an ASI is and can often result in confusion. Further, the large number of conflicting terms used in media and in government can lead to misunderstanding and mislabeling, which affects data collection and therefore research. Clarification is needed and this research includes consideration of the published terms and definitions with the addition of amplifying information.

First, an active shooter is thought of by the FBI (2021, p. 1) as "… one or more individuals actively engaged in killing or attempting to kill people in a populated area. Implicit in this definition is the use of one or more firearms.". This is a fairly straightforward explanation of what an active shooter is. Therefore, an ASI is an incident where an active shooter commits his crime. One might think that this is all the reader needs to understand the verbiage of active shooting incidents. However, some details require recognition to better understand the research conducted in this work and to have a clearer view of what specific properties an ASI contains. It is important to recognize, as Krouse and Richardson (2015, pp. 2-3) point out, that "… statute, media outlets, gun control and rights advocates, law enforcement agencies, and researchers often adopt different definitions of 'mass killing,' 'mass murder,' and 'mass shooting,' contributing to a welter of claims and counter-claims about the prevalence and deadliness of mass shootings.". This research is not intended to muddy the waters further, only to clarify a more refined definition of an ASI.

As Smart and Schell (2021) note:

> There is no standard definition of what constitutes a mass shooting, and different data sources—such as media outlets, academic researchers, and law enforcement agencies—frequently use different definitions when discussing and analyzing mass shootings. For instance, when various organizations measure and report on mass shootings, the criteria they use in counting such events might differ by the minimum threshold for the number of victims, whether the victim count includes those who were not fatally injured, where the shooting occurred, whether the shooting occurred in connection to another crime, and the relationship between the shooter and the victims. These inconsistencies lead to different assessments of how frequently mass shootings occur and whether they are more common now than they were a decade or two ago. (p. 1)

So, according to at least one source, the term "mass shooting" does not appear to be well defined in any official capacity (Smart & Schell, 2021). However, one can easily picture what the term is and what it means. A "mass casualty shooting" is mentioned by the United States Department of Justice for Victims of Crimes as the murder of three or more individuals (United States Department of Justice Office for Victims of Crimes, 2017). This mimics the FBI's definition of "mass killings", which is defined as three or more deaths (Krouse & Richardson, 2015). These terms are also very close to the definition of "mass murder" which is multiple homicide incidents in which four or more victims are murdered in one event and in one or more locations (Krouse & Richardson, 2015). One can see that these terms are evolutionary in the field of criminology and have their history in law enforcement investigations over decades (Krouse & Richardson, 2015). Another definition that can help understand the topic is "mass casualty incident". DeNolf and Kahwaji (2020, p. 1) give an apt explanation of it as "an event that overwhelms the local healthcare system, where the number of casualties vastly exceeds the local resources and capabilities in a short period of time.". Clearly, many factors play into these definitions, and according to current and former FBI agents, these crimes can be classified by victim counts, type, and style (Krouse & Richardson, 2015).

Knowing this information, the author can better define an "active shooting incident". First, an active shooting incident is not related to any other crime being perpetrated immediately before, during, or after the incident. For example, if violent gang members attack each other and three or more people are killed, this is not an ASI. Further, should a person kill four or more of their family members, even in an open and public space, it should not be counted as an ASI. The FBI defines this second type of killings as "Familicide" (Krouse & Richardson, 2015). Secondly, an ASI does not have to produce casualties. What matters in the definition is not the number of innocents killed or wounded but that one or more shooters intended to inflict such casualties. This is easily determined by their actions, such as how many rounds they fired, whom they appear to target, their various tactical movements, and later, what investigative measures are conducted. However, firearms must be discharged to classify such an occurrence as an ASI since the "active" property is important. Thirdly, the type of targets a shooter selects must be what one could refer to as "soft targets". For example, unarmed civilians of any age are soft targets, as opposed to armed police officers. Though a shooter attacking a police station might be an ASI in some minds, the author would consider this a terrorist attack. This is because the police represent authority and law; hence the shooter's intent is clearly to affect politics in some manner. Further, given the generally accepted understanding of the mental state of an active shooter, law enforcement is generally not in the target set of these people. Most shootings against police that ended in more than four deceased officers were clearly motivated by politics, such as the Dallas shooting (Okoro, 2021) (Associated Press, 2016). Other shootings that produce deaths or casualties among law enforcement are related to other crimes, making the definition of ASI invalid. "Intent" of a criminal is often hard to measure and mostly irrelevant to law, yet to narrow the collection of data that researchers and policy makers rely on, intent should be considered for ASIs.

Lastly, the victims of the shooter could have a certain randomness in target and location selection. That is, the shooter intends to injure or kill soft targets yet might also engage the police.

Also, they are not necessarily particular in the location they choose. They might also decide to attack their colleagues, friends, or family while engaged in their act. Some school shootings are examples of this, where a shooter targeted students, faculty, and staff and, in the case of Columbine, also shot at police. This is different from a disgruntled employee who specifically targets certain coworkers for murder. In the case of a disgruntled employee, the author would consider that instance a mass murder or mass killing, and not in the same category as an ASI.

The author's definition closely resembles that of Schildkraut and Elsass (2021):

> A mass shooting is an incident of targeted violence carried out by one or more shooters at one or more public or populated locations. Multiple victims (both injuries and fatalities) are associated with the attack, and both the victims and location(s) are chosen either at random or for their symbolic value. The event occurs within a single 24-hour period, though most attacks typically last only a few minutes. The motivation of the shooting must not correlate with gang violence or targeted militant or terroristic activity. (p. 1)

The researcher understands and welcomes any criticism of the author's definition. There are still a few cases where various properties overlap, or the instance cannot be declared an ASI exclusively. Nevertheless, it is hoped that at least for this work, the definition helps in illuminating the research.

2.3 Active Shooting Incident Statistics

To provide the reader with more context and information, the author included some of the most important active shooter statistics in this work. Though some of this data is cast throughout this paper, it was judged beneficial to provide all of it within its own section for convenience and thoroughness.

According to the FBI, there have been a total of 333 active shooter incidents between 2000 and 2019, and 135 of those had three or more people killed, meeting the "mass killing" definition (Federal Bureau of Investigation, 2021). A total of 2,851 casualties occurred, excluding the shooters, and 1,062 victims were killed while 1,789 were wounded (Federal Bureau of Investigation, 2021). Of all casualties, 95.6% were civilian, the others law enforcement or security guards. Of the 345 shooters, 332 were male and 13 were female. Ten incidents involved more than one shooter (Federal Bureau of Investigation, 2021). Another study, looking at 314 public mass shootings in the USA from 1966 to 2016, suggests that the average age of a shooter is 35 (Silva & Capellan, 2019). Another article puts the average age over a similar time period, including shootings up to 2020, at 33.2 (Jaclyn Schildkraut & Elsass, 2021). More than 50 percent of shooters are aged 30 or above (54%) (Schildkraut et al., 2018). This invalidates the common misconception most active shooters are teenagers or individuals in their 20s.

Out of all shooters listed in the recent literature, 119 committed suicide, while police arrested 150, and 71 were either killed by police or civilians (Federal Bureau of Investigation, 2021). Only five of the shooters remain at large at the date of publication of the report. Out of the 12 location types identified by the FBI, the majority of the shootings occurred at businesses open to the public (96), followed by open spaces (50), then Pre-K-12 schools (44), and businesses closed to the public (41) (Federal Bureau of Investigation, 2021). One can tell that the majority (137) of these shootings occur in businesses of various types. The reference clarifies that the number of incidences has increased from the year 2000, given that the year 2000 has three events, and the year 2019 had a total of 30 (Federal Bureau of Investigation, 2021). Figure 2.1 below demonstrates the increase:

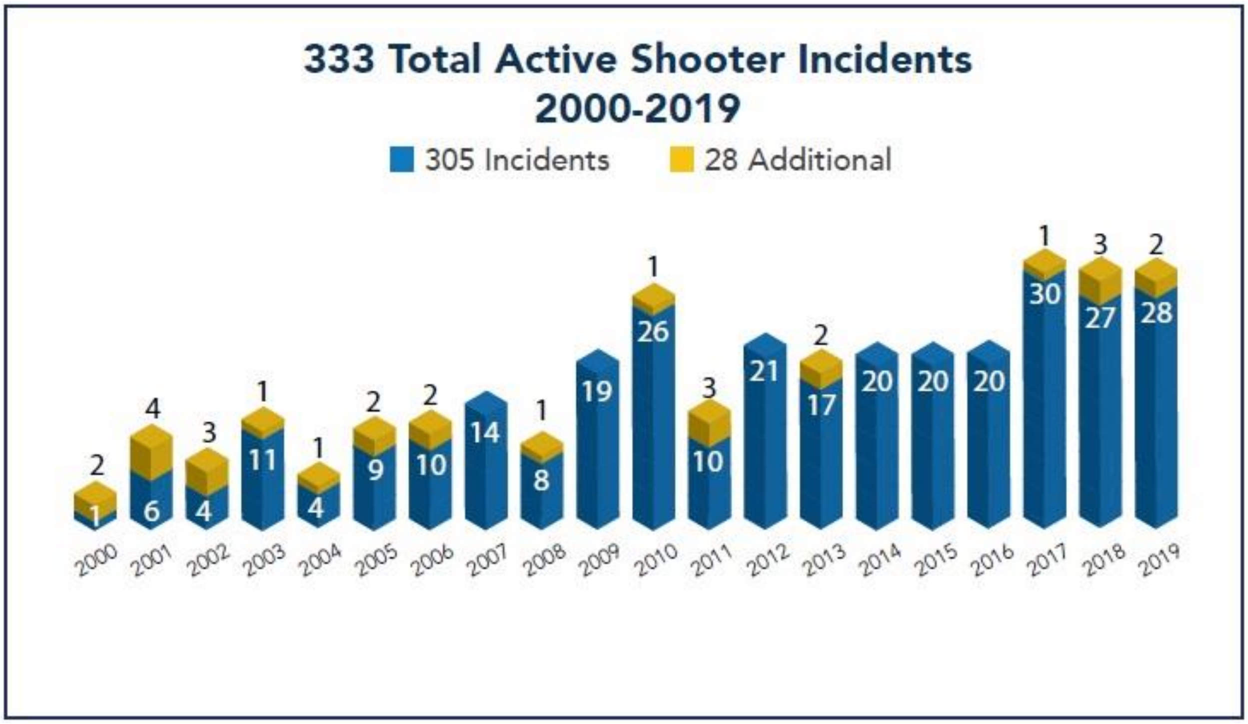

Figure 2.1 - Total Active Shooter Incidents 2000-2019
(Federal Bureau of Investigation, 2021)

As far as the weapon types used by shooters, the FBI (2021, p. 30) states that: "In the 333 active shooter incidents, handguns accounted for 67% of the weapons used, 38% of the 345 shooters had multiple weapons, 5% wore body armor, and 4% had access to or deployed additional devices.". Additionally, in only 26% of the incidents, shooters used long guns (rifles of any type), and 10% used shotguns (Federal Bureau of Investigation, 2021). Blair and Schweit (2014, p.8) found that out of 160 active shooting incidents analyzed, "63 incidents where the duration of the incident could be ascertained, 44 (69.8%) of 63 incidents ended in 5 minutes or less, with 23 ending in 2 minutes or less.". As already discussed, most of the shootings end in suicide of the shooter or via force used by police or others to stop it. Police response times are critical, given the short duration of the incidents and their lethality. One article highlights "The Stopwatch of Death" factor. This references the Virginia Tech shooting where there was nearly eight murders or attempted murders per minute (Police1, 2007). Police response times vary across the nation and are dependent on a variety of factors. The median response time to active shooting incidents is three minutes (Jonson et al., 2020). It is not clear from the reference if this response time is "victim time" or "dispatch time" though if it is the latter, it would mean that the shooter has even more than three minutes to cause casualties.

2.4 Active Shooting Incidents and the Evolution of Responder Training and Tactics

Active shooting incidents are not a new or modern phenomenon. They have existed in one form or another since the early days of gunpowder weapons. If one applies some of the modern definitions of "active shooting incident" to the past, one can see that violence with firearms against civilians has been ripe throughout history.

On August 1st, 1966, a shooter ascended the clock tower at the University of Texas with multiple weapons and began a shooting that would rank in importance second only to the war in Vietnam in the minds of Americans that year (Colloff, 2006). In that incident, 43 people were shot, and 13 of those perished at the shooter's hands (Colloff, 2006). According to Colloff (2006, p.1), the shooter "introduced the nation to the idea of mass murder in a public space". This incident also prompted authorities all across the nation to develop Special Weapons and Tactics (SWAT) teams in order to have trained responders for such events (Colloff, 2006). Further, Colloff (2006) notes several other important actions and responses during and after the incident:

> Students waited and waited for the police to arrive. The shootings would spur the creation of SWAT teams across the country, but at that time, the Austin Police Department had no tactical unit to deploy. Its officers had only service revolvers and shotguns, which were useless against a sniper whose perch was hundreds of yards away. Communication with headquarters was difficult, with few handheld radios, and the phone system was jammed across the city. Some officers went home to get their rifles; others directed traffic away from campus. In the absence of any visible police presence, students decided to defend themselves. (p. 17)

Delving a bit deeper into history, the first of what modern Americans would call an active shooting occurred on September 6^{th}, 1949, though this incident was simply called what it was at the time, mass murder (Sauer, 2015). It was perpetrated by a deranged individual using a firearm and ended the lives of 13 people, including children. Even generations ago, both civilians and responders were ill-prepared for the kind of carnage a shooter of this magnitude could unleash on his community. It took the shooter 20 minutes to conduct his actions, including the time it took police responders to apprehend him after a gun battle and the use of tear gas to force a surrender (Sauer, 2015). This event was so shocking and unbelievable, several of the residences of Camden, NJ, demanded the murderer be lynched on the spot, and one of the police officers asked the shooter, "What's the matter with you? You a psycho?" (Sauer, 2015, p.10). Perhaps, this is an early indication of the state of mind of such a vile individual and their general mental health, something that one study addresses in some detail and that is worthy of note regarding other mass casualty shootings throughout history (Frantz, 2021). This event is another reminder that these types of violent acts have occurred well before the mass murder at Columbine High School. All nations and societies are subject to violent acts by individuals and groups of individuals who are motivated to do so for various reasons.

It is a well-understood fact that active shooting incidents (ASIs) have been increasing steadily in the frequency of occurrence since the late 1990s (Federal Bureau of Investigation, 2021). Some research also shows an increase in shootings since the mid-1960s (Jaclyn Schildkraut & Elsass, 2021).

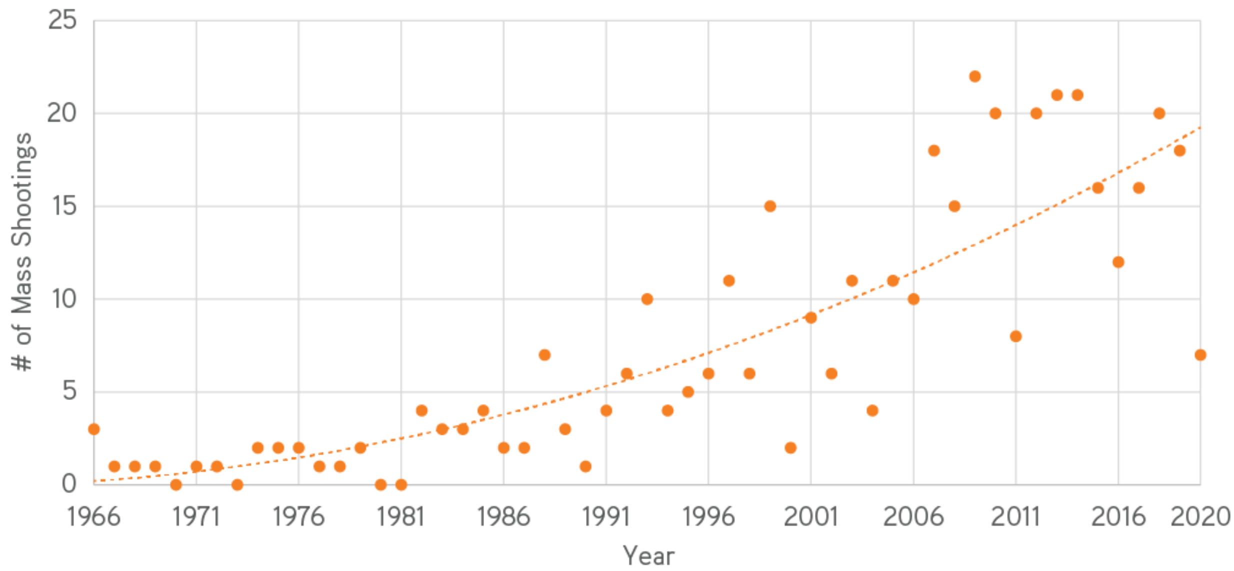

Figure 2.2 - Mass Shootings 1966-2000
(Jaclyn Schildkraut & Elsass, 2021)

The most noteworthy event that started what we might call the "modern era of active shooting incidents" was the Columbine High School shooting on April 20th, 1999. This evil was perpetrated by two students at the high school and resulted in 13 people being murdered and many more injured (Editors, n.d.). Much of the reason for so many casualties were the apparent lack of preparedness and understanding of the event unfolding – in both civilians and responders. Neither the teachers nor the students acted in time to either run or hide from the shooters or, as a last resort, fight them. Further, the police responders failed to immediately enter the school in force and engage the shooters, given that this was not the local or nationally accepted tactic at the time (Blair et al., 2013). The after-action review of this shooting created the impetus to significantly change responder tactics in a relatively short time, and create programs to help train civilians to act better in a deadly situation (Blair et al., 2013). As Philips (Phillips, 2020) points out:

> An active shooter event presumes an "expeditious" resolution, which decreases the safety benefits associated with a slow and deliberate response. Police training has also focused on the role of the officer to immediately engage a shooter, demonstrated in popular policing-related websites and training organizations. (p. 266)

Today's law enforcement first responders are trained and encouraged to bring an active shooting to a swift end, and over the last two decades many police departments and other agencies have received detailed and effective training to do so (Blair & Martaindale, 2019). These changes in tactics, techniques, and procedures were apparent on April 16th, 2007, when an ASI occurred on the Virginia Tech campus. Initially, this shooting began with two victims and a murder investigation by local law enforcement (Virginia Tech Review Panel, 2007). This response, though not ideal for ASIs, was largely the correct one given the information available in the early hours of a protracted event. The shooter had attacked two students near his residence hall, then delayed several hours to continue his murder spree in a different building across campus (Virginia Tech Review Panel, 2007).

The primary events that produced the majority of casualties occurred in Norris Hall at 9:40 am and lasted until 9:51 am. Officers responded within three minutes of the shooting's start, and after determining where the shooting was ongoing, attempted to enter the building. However, many of the doors were chained shut from the inside by the shooter (Virginia Tech Review Panel, 2007). The procedure of immediately entering the building is the direct result of better training and education of first responders, based on the recognition that shooters will not negotiate and seek only to kill as many innocents as possible (Blair et al., 2013) (Blair & Martaindale, 2019). As certain research also suggests, closing with and engaging the shooter sooner rather than later will increase the survival rate of all civilians caught in the incident (Bott et al., n.d.). The Virginia Tech shooting is arguably the first high-profile shooting after Columbine that proved the new tactics correct. The shooter in this situation overheard police officers using a shotgun to break a small lock to gain entrance then proceeded to shoot himself, effectively ending the incident (Virginia Tech Review Panel, 2007). Based on historical data, it is certain that had police used the same tactics, techniques, and procedures that were employed at Columbine, it would have resulted in far more casualties.

Another tragedy the researcher would be remiss to overlook is the school shooting on December 14th, 2012, in Newtown, CT (Ray, 2020). Also known as the "Sandy Hook Elementary School Shooting", this incident was particularly troubling to the nation due to the nature of the victims. The perpetrator ended the life of 27 individuals, including seven adults and 20 children, before committing suicide himself (Stephen J. Sedensky III, 2013). Sedensky (2013) lays out the timeline of the event that tells us how quickly life is lost and how important it is to have responders on the scene as fast as possible:

> The response to these crimes began unfolding at 9:35:39 a.m. when the first 911 call was received by the Newtown Police Department. With the receipt of that call, the dispatching and the arrival of the police, the law enforcement response to the shootings began. It was fewer than four minutes from the time the first 911 call was received until the first police officer arrived at the school. It was fewer than five minutes from the first 911 call, and one minute after the arrival of the first officer, that the shooter killed himself. It was fewer than six minutes from the time the first police officer arrived on SHES property to the time the first police officer entered the school building. In fewer than 11 minutes twenty first-grade pupils and six adults had lost their lives.

Within eleven minutes, the shooter was able to bring havoc and death to over two-dozen people, the majority children. This emphasizes how little time civilians have to react to an incident like this and how few precious minutes police have to respond. If it was not for the immediate response of the police, more innocent lives might have been lost. As research suggests, more lives could have been saved had there been responders on-site, such as a school resource officer (SRO) or school faculty with a concealed carry weapon (CCW) (A. Kirby et al., 2016) (Bott et al., n.d.). Further, the hardening of various doors, including the outer glass door, could have prevented some or perhaps all of the deaths, as other studies indicate (A. Kirby et al., 2016). Yet, the fact that police officers responded quickly and immediately entered the building shows the successful evolution of tactics, techniques, and procedures since the Columbine massacre. However, this incident still demonstrated the need for better training of civilians, the hardening of facilities, and the increased need for armed protectors on site.

A recent incident, as of the writing of this dissertation, highlights the improvement of police training and tactics, and brings into focus the courage officers must demonstrate to apply them effectively. It further exposes the error of losing speed, surprise, and violence of action by responders in the early stages of an incident and how this potentially costs the lives of innocent civilians. On March 22, 2021, at approximately 2:30 pm a shooter began his murder spree in the parking lot of the King Sooper store in Boulder, Colorado (Larsen, 2021) (Bradbury, 2021). At 2:35 pm three officers arrived on the scene and entered the building within 30 seconds (Jennifer Campbell-Hicks, 2021). As a video that was live-streamed during the event demonstrates, the police officers entered the store and were immediately engaged by the shooter (ZFG Videography, 2021). Sadly, one police officer was fatally wounded by the shooter (McBride, 2021). The actions of this officer and his two colleagues were vital in providing resistance to the murderer and arguably helped save more lives. Speed, surprise, and violence of action are required when engaging a shooter bent on killing, and modern training and tactics reflect this truth (Blair & Martaindale, 2019) (Texas State University, 2021).

Unfortunately, instead of pushing ahead and eliminating the threat, despite their tragic casualty, the other officers stalled, withdrew from the shooter and took cover at the entrance to the store (ZFG Videography, 2021). This created a standoff, something that modern policing knows is the exact opposite of what needs to happen in an active shooting incident. A standoff delays medical care to wounded individuals, increasing the possibility of more avoidable deaths. In fact, the United States military refers to the first hour after being wounded as the “Golden Hour”.

Decades of experience in war has shown that survival rates plummet one hour after being wounded (Rasmussen et al., 2015). This, of course, also applies to anyone injured by weapons that are similar to those found in modern war. When accounting for triage and evacuation, a victim might only have 30-45 minutes to receive medical care, so it is incumbent upon the responders to end a shooting as quickly as possible, even if it means putting their bodies and lives at risk. In the end, the mass murderer at Boulder ended the lives of ten people, including one police officer (Jennifer Campbell-Hicks, 2021). One wonders how many wounded might have been saved if the other officers had acted and eliminated the threat quickly, allowing for expedited treatment and evacuation of the wounded.

The final active shooting incident to highlight in this section is the tragic event that transpired in Parkland, Florida. On February 14th, 2018, at around 2:20 pm a gunman entered building 1200 at Marjory Stoneman Douglass High School, intent on murder (Alanez et al., 2018). Within 5 minutes and 32 seconds from first to last shot fired, 17 people were murdered and 17 wounded (Alanez et al., 2018) (Editors, 2019). The gunman, armed with an AR-15 rifle and equipped with tactical gear, was able to conduct his heinous acts unhindered by any resistance on three separate floors of the building (Alanez et al., 2018) (Broward Sheriff's Office, 2018). This tragedy is made even worse, given the fact that nearly 20 years of proper first responder tactics, techniques, and procedures were either not trained to or not carried out by Broward County Sheriff's deputies and others (O'Matz et al., 2018). This is in stark contrast to the Coral Springs officers, who were properly trained to respond to an active shooter and performed admirably (O'Matz et al., 2018). The people who could have affected the situation positively, specifically the SRO and other school administrators and staff, failed in their assigned duties from the outset (Alanez et al., 2018).

As mentioned above, police have been trained since Columbine to charge towards the sound of gunfire in hopes to interdict a shooter as soon as possible, which would likely end the shooting and therefore save lives. This was not the course of action chosen by the SRO that day. Instead, he hid for the duration of the shooting outside of the building (Alanez et al., 2018) (MSD Public Safety Commission, 2019, pp. 13-14). This was an unconscionable action by an SRO sworn to protect students against such violence. Some people argue not to judge the SRO too harshly, given the confusing situation between the sound of fireworks and gunfire (Joel F. Shults, 2019). However, looking at the history of training responders and all other evidence discussed, this is a questionable conclusion, and the State moved forward in charging this individual for multiple crimes (MSD Public Safety Commission, 2019, p. 14). A further breakdown in procedure occurred even before the SRO was involved. A campus watchman failed to alert the school via radio of a “code red”, an order to lock down the school, when he witnessed the shooter entering building 1200 with what appeared to be a rifle bag (Alanez et al., 2018). Regarding the watchman, Alanez et al. (2018) states:

> He recognizes Cruz as "Crazy Boy," the former student that he and his colleagues had predicted most likely to shoot up the school. He radios another campus monitor/coach, but he does not pursue Cruz and does not call a Code Red to lock down the school.

Clearly, if either or both individuals had acted according to the proper procedure and their duties, things might have been different that day. Even the responding officers from the Broward County Sherriff’s department failed as a whole to respond in the most advantageous manner and other school staff also performed in a poor manner. One was a monitor in the halls who saw the shooter in the hallway on the first floor of building 1200 before the murderer had pulled his gun from the bag (Alanez et al., 2018). The monitor immediately ran away, failing to alert the school administration of the danger (Alanez et al., 2018). What makes the actions of this individual even worse, as soon as he heard gunshots, he ran up the stairs and hid in a janitor's closet, failing to raise the alarm immediately, even as he had a radio to do so (Alanez et al., 2018) (MSD Public Safety Commission, 2019, p. 233).

Another school staffer was informed by a student that there was someone on campus with a weapon (Alanez et al., 2018). The staffer did not call a "code red" either. Despite all the mistakes, and poor actions of some, a few performed their duties to the best of their ability. Several staffers ran towards the gunfire in building 1200, and some paid with their lives for their actions (Alanez et al., 2018). At least one teacher on the third floor took charge and directed students out of the hallway and into classrooms as the shooter was still engaged on the second floor (MSD Public Safety Commission, 2019a).

The Coral Springs police officers responded as trained. They used speed, surprise, and violence of action, intent on eliminating the threat even while acting with minimal information (Alanez et al., 2018) (MSD Public Safety Commission, 2019a, p. 34). Again, these are the correct tactics to use for an active shooting incident as a responder, ever since the massacre of Columbine changed the view of such events. Fortunately, appropriate, and life-saving responses came from some of the students and staff themselves. It is important to point out the correct actions taken by students and teachers on the second floor of building 1200. By concealing themselves in the classrooms, none of them were shot (MSD Public Safety Commission, 2019, p. 27). Hiding or avoiding the shooter reflects one of the tactics civilians can use in an active shooting incident to increase their chances of survival.

It is important to note the good and bad actions taken by various individuals in response to a shooter so lessons may be learned. The historical events discussed in this section have been some of the most publicized in recent history. They can instruct on implementing policies that might increase the chances of survival for anyone caught in an ASI.

2.4.1 Open-Air Venues and Open Spaces

When most people imagine an active shooting, they picture it occurring in a closed space such as an office building, a school, or another place where individuals gather daily. This is not a false image, given the data on the subject. As the FBI notes, the majority of active shooting incidents occur in some form of closed space (Federal Bureau of Investigation, 2021). However, a notable number of active shooting incidents occur in what the FBI defines as open spaces. The agency states that out of 333 active shooting incidents, 50 occurred in open spaces (Federal Bureau of Investigation, 2021). One can infer from the reference that some spaces, such as malls, might also have various open spaces, even if they are in a separate category of the 12 listed by the FBI. It is vital to note that a significant number of these shootings occur in open spaces to warrant attention. A number of high-profile and well-publicized shootings in open spaces occurred in the last few years within the United States.

The deadliest and most casualty-producing active shooting incident in American history was the Las Vegas Route 91 Harvest music festival shooting (Corcoran et al., 2019). On October 1st, 2017, a shooter engaged approximately 22,000 people from a hotel window adjacent to an open-air concert venue (Joseph Lombardo, 2018). The shooting itself lasted about ten minutes, during which the murderer injured 413 individuals and killed 58. Additionally, 456 people suffered injuries related to the shooting but were not related to gunshots or shrapnel (Joseph Lombardo, 2018). This shooting was notable not only in the number of casualties produced but also in the type and number of weapons the shooter used. The shooter was shown to have had 14 separate weapons in the hotel room, including several AR-15 and AR-10 type weapon systems (Joseph

Lombardo, 2018). After creating such carnage, the murderer expended over 1000 rounds during his spree from those 14 weapons while ending his life with a Smith and Wesson revolver (Joseph Lombardo, 2018). Also of note was a device the shooter used to mimic full-automatic fire with the otherwise semi-automatic weapon systems. He used what is commonly referred to as a "bump stock", which uses the recoil of the weapon into one's shoulder to allow for the firing cycle to repeat in a quasi-automatic mode (ATF, n.d.). This, of course, contributed greatly to the lethality of the attack, yet would not have been possible had the venue not been an open space with tens of thousands of people in attendance. This highlights the special dangers associated with open venues and spaces, particularly large events such as this.

The police response to the shooting was immediate and effective. This was largely due to the security presence in and around the event, and also due to the training received by officers (Joseph Lombardo, 2018). The quick response demonstrates the wisdom of having armed and trained responders on-site to engage a shooter as soon as possible to avoid more death and injuries. One can argue that the police were not a key factor in stopping the shooter since he committed suicide long before the special weapons and tactics (SWAT) team entered the room (Corcoran et al., 2019) (Joseph Lombardo, 2018). However, given the tactics they are trained to use, one can also make the case that they would have had a higher sense of urgency and would have acted faster and more aggressively if the shooter had not stopped shooting after ten minutes. The scale of the venue and the number of people present demonstrate the difficulties responders have in coordinating an immediate and effective response to a shooting as well as its aftermath (Marcou, 2019). This is just one factor that makes open-air venues more vulnerable and harder to secure. Other factors that can impede a quick and effective response are communications issues, lack of incident command training, and proper planning to name just a few (Marcou, 2019). All of these are factors in any active shooting incident, yet they clearly are magnified in open spaces and large open-air venues.

The final shooting that occurred in an open-air venue which the author wants to discuss is the Gilroy Garlic Festival shooting. One purpose of highlighting this shooting is that it is instrumental in the experimentation the author performed for this work. It would be unwise not to delve into the details of this shooting.

The Gilroy Garlic Festival is held annually and attracts up to 100,000 people over three days (Frantz, 2021). On July 28th, 2019, at 5:51 pm, a shooter entered the festival grounds from the north after breaking through a locked gate with bolt cutters (Rosen, 2020). The shooter proceeded to shoot 20 individuals and kill three within about one minute (Johnson et al., 2019). The shooter was equipped with an AK-47 style semi-automatic rifle, several 30-round magazines, and one 75-round "drum" magazine. He was also equipped with a protective vest, capable of stopping rounds fired from police handguns (Rosen, 2020).

The shooter began to engage civilians on the festival grounds while stationed near a blow-up slide on the northwest corner of the grounds. He experienced a weapons malfunction and struggled to clear it for 15-20 seconds before changing magazines and resuming his murder spree (Rosen, 2020). Unfortunately, even after getting off one shot into the ground before the malfunction, this initial round was not enough to trigger a reaction by responders or even most of the witnesses on the grounds (Rosen, 2020). However, as soon as the shooter started firing again, police officers approximately 150 meters south of the shooter began to react as their training and instincts took over. A veteran of campaigns in Iraq and Afghanistan, one officer recognized the distinctive sound of enemy weapons fire and immediately charged towards the gunfire (Rosen, 2020).

The officers were able to engage the shooter quickly, forcing a withdrawal and ultimately suicide, ending the incident. This event is a prime example of how the immediate actions of first responders and their proper pre-positioning saved lives. Also of note is how and where civilians fled once they understood the threat. According to one video posted online, individuals were fleeing away from the sound of gunfire in confusion and panic (NBC New York, 2019). It is well known by law enforcement and personal security instructors that most people will flee in the direction of the entrance they came through, be it a closed space or an open space. However, one must wonder what these people would do if the only known exit were in the direction of the threat. This is a problem that can plague any venue, especially one with a large number of people present. Further, civilians running away from a shooter might hinder first responders from properly engaging the threat.

Other researchers have written about the differences and the increased susceptibility of open spaces to active shooting incidents. Regarding open spaces and venues, Frantz (2021) states:

> Many open-air venues are not in place permanently, setting up for short periods of the year to host events and creating a unique problem set for event planners, who may not have their entire staff until the actual event day(s). Time is not available for proper planning by event management in these situations and no standard active shooter protocol exists for rapid implementation across all variations of open-air venues. (p. 23)

This emphasizes why open-air venues and open spaces, in general, are more vulnerable to active shooters. The physical size of the environment and the number of people in the area also contribute to the vulnerability. We know that the shooter intends to cause as much harm as possible (Frantz, 2021). Frantz (2021) determined the following regarding active shooters:

> There is no definitive method to predict the actions of an active shooter. Individuals that choose to carry out a mass shooting are mentally unstable and follow an unpredictable path. The majority of active shooters are also suicidal (Knox, 2018) which can push a person to the edge, with no regard for harm to themselves. Mental instability creates actions that are near impossible to predict, but identification of behavioral patterns can lead to better understanding of how a shooter will handle

> situations that lead to them using a firearm against innocent victims. Given this simple fact, researchers should pay particular attention to open-air venues and open spaces and find ways to mitigate damage specifically to these environments. (p. 35)

If the shooter is mentally unstable, seeking to cause maximum carnage during an event, then open spaces and open-air venues are optimal targets. As seen with the Las Vegas shooting in 2018, there were plenty of opportunities to take innocent life in a very short period of time. It is the author's hope that more attention is paid to these venues by law enforcement and other researchers in the field.

2.5 Agent-Based Modeling and Simulation

Agent-Based Modeling and Simulation (ABMS) has become a widely used and highly effective tool for researchers in government and private organizations. The author will focus on ABMS in the context of computerized modeling and simulation. Bandini et. al. (2009, p.1) state that a "… computer simulation is related to the usage of a computational model in order to improve the understanding of a system's behavior and/or to evaluate strategies for its operation, in explanatory or predictive schemes". In broader terms, a computer simulation means that we wish to gain insight into a system's behavior using a computational model (Bandini et al., 2009). We simulate a real-world scenario, event, or environment to observe the behavior of its components and extract information. Computer modeling and simulation is related to manipulating a computational model for analysis of a system and assessing strategies for prediction and description (Abar et al., 2017). A model is an abstract and simplified representation of reality, existing or planned, and these models are often defined to explain existing phenomena or predict future phenomena (Bandini et al., 2009). Exploratory modeling is a method that can provide value to researchers even if a model cannot be validated and fails to predict anything of immediate value (Bankes, 1992). This is true, given that we are searching for insight from a model, not necessarily raw numbers, and even partial information can assist policy makers in making decisions (Bankes, 1992). Often, a model can help reduce uncertainty even if the results are questionable and validation is impossible. One must consider these limitations when presenting the results of such a model. Bankes (1992, p.12) states that in exploratory modeling, "… the computer functions as a prosthesis for the imagination, allowing the discovery of novel explanations of known facts or unrealized properties …". One of the most important benefits of using ABMS, or any simulation for that matter, is that it allows researchers to conduct their studies without risking real lives or property (Bandini et al., 2009). One can model a scenario of past incidents or a present situation and, depending on the abstraction and fidelity of the model, gain insight into what might have been done differently in the past or what could be done differently in the future. This is the key benefit of any modeling, and in particular computerized modeling, since modern computing systems provide us enormous processing power, storage space, and the ability to visualize data.

Agent-based modeling and simulation adds to this power by better reflecting the world as we understand it since it maps well to most situations. ABMS is a bottom-up approach that lends itself better to the way most humans analyze a problem. Within an agent-based model, an agent is considered to be an independent component that acts within the environment and interacts with it and other agents (Macal & North, 2005). Macal and North (2005) state that "The fundamental feature of an agent is the capability of the component to make independent decisions. This requires agents to be active rather than purely passive.". Some would consider an agent to be similar in concept to an object within object-oriented software. A well-defined entity that is instantiated in computer memory into a known state and that has its own encapsulated data and abstracted behavior, which can also interact with other parts of the software through a well-defined interface.

An agent can also be autonomous, self and goal-directed and might even show aspects of being able to learn within the system (Macal & North, 2005).

In some cases, an agent can model an entire system such as a large corporation or an entire country; in other situations, a human being, or an inanimate object like a traffic light. This depends on the simulated environment, the researcher's model, and what he or she wishes to gain from it. In recent years, some have increasingly used Machine Learning (ML), and specifically Deep Reinforcement Learning (DRL), to give agents capabilities yet unseen in ABMS. For example, DRL was used together with agents to model charging loads for electrical taxis in a more efficient manner (Jiang et al., 2018), to handle electricity market data to maximize returns (Kiran & Chandrakala, 2020), and to enhance cyber security by using DRL agents to defend against Distributed Denial of Service (DDoS) attacks in a network (Xia et al., 2019).

To properly implement models efficiently and effectively, the right tools are needed. Abar et. al. (2017) list a total of 84 toolkits available for professionals to use to conduct model development and research. Included in that list is AnyLogic (AnyLogic, n.d.), the preferred ABMS tool of the Purdue Homeland Security Institute (PHSI), where ASI research is conducted to inform policy. ABMS is an effective methodology to allow various researchers and other professionals to explain systems and inform policy. Exploratory modeling can be useful to researchers even when the models cannot be validated, though extreme caution must be taken in those cases. The primary value of any computerized modeling, and ABMS specifically, is that it allows us to experiment in virtual situations and adjust variables without risking life or property. If a model can be properly validated, the possibilities to test certain variables to observe outcomes is nearly endless.

Before continuing, it is vital to define one of the most critical concepts in ABMS, and modeling in general, namely fidelity. The fidelity of a model is fundamental given that it distinguishes what makes up a computer simulation compared to any other software program (Gross, 1999). Gross (1999, p. 1) states that "… the unique measure of goodness for simulation is how well the simulation makes its representation, or its ***fidelity***" and defines it as "The degree to which a model or simulation reproduces the state and behavior of a real world object or the perception of a real world object, feature, condition, or chosen standard in a measurable or perceivable manner;" (p. 2).

2.5.1 Active Shooting Incident Research using Agent-Based Modeling and Simulation

Much work has been done in the last two decades within the active shooting incident (ASI) research field. One of the main reasons for increased ASI research was the Columbine High School shooting in April 1999. This tragedy sent shockwaves through the nation and preceded two decades of increasingly high-profile shootings in schools, businesses, and public places. These shootings have produced a significant amount of research into the topic of ASIs, as well as spawned a professional class and programs that deal with school shootings specifically (ALICE Training Solutions, n.d.) (City of Houston, 2012) (National Rifle Association of America, n.d.). In support of the professional trainers and the policy makers is academic research. In this important field, ABMS has assisted greatly in determining the best policies to safeguard life and property by recommending concrete actions.

For example, Kirby et.al. (2016, p. 6) determined that "… the simple policy of locking doors can reduce the number of people shot in an active shooter event by almost 25%.". This study further determined that concealed carry at the workplace can also significantly reduce overall casualties, given the immediate response to a threat (A. Kirby et al., 2016). Another study that used ABMS looked at ways to mitigate the impact of active shooting incidents in schools by adjusting various policy options and their effects (Anklam III et al., 2015). The study used the AnyLogic system to produce an agent-based model to analyze various scenarios. Anklam III. et.al. (2015, p. 11) discovered that to "… decrease the number of casualties, the response time must be reduced." and that the most efficient way to do that is to have armed personnel present at the school. This could be accomplished by having armed School Resource Officers (SROs) and armed administrators and teachers present before the incident (Anklam III et al., 2015).

A significant contribution to ASI research using ABMS was conducted by modeling the 1999 Columbine event (Lee, 2019). This model re-created the Columbine High School library with high fidelity and relatively low abstraction to ultimately test the run, hide, fight methodology (Lee, 2019). It validated the model by recreating the event in detail and running a scenario that produced similar results as the historical event. It also extracted important information on the historical event such as shooter weapon discharge rate, shooter movement, and other important variables (Lee, 2019). These variables were then used to test the run, hide, fight methodology and its effectiveness within the historical environment of the Columbine High School (Lee, 2019). Also, this established a baseline that allowed follow-on researchers to compare against. This work was able to determine that RUN.HIDE.FIGHT® is effective at reducing casualties, and that offensive measures are the most effective.

Expanding on this agent-based model, another team of researchers used it to add an SRO and a concealed carry weapon (CCW) holder to the library (Bott et al., n.d.). These additions contributed to the offensive capabilities of the civilians within a short period of time after the incident started. The results of adding an SRO and a CCW holder to the library were to reduce casualties significantly compared to the historical incident or to running or hiding (Bott et al., n.d.). In this case, ABMS for ASI has produced policy suggestions that, when implemented, can increase survival during an incident. A very important dissertation using ABMS and concerning ASI determined that many different policies such as locked doors, concealed carry, and rapid response all lead to reduced casualties (A. M. Kirby, 2016). In this piece, ABMS was used to model various scenarios to reach conclusions on the recommended policies (A. M. Kirby, 2016). It was determined that simply locking doors was most effective, given the default of not having any defense, and that providing offensive capabilities further reduced casualties.

Of note is one study that looked at the potential effectiveness of the proposed assault weapons and high capacity magazine bill of 2013 (Hayes & Hayes, 2014). It re-created the historical event of the Aurora Colorado theater shooting and validated the model with this scenario. It tested the stated hypothesis by running indoor and outdoor scenarios. The researchers were able to determine, among other things, that the bill as proposed would have had negligible effects on the number of casualties incurred (Hayes & Hayes, 2014). It also determined that the weapons' rate of fire was the primary driver of the number of people shot during the simulated event (Hayes & Hayes, 2014). Another noteworthy agent-based model looked into how unarmed individuals would perform when tasked with subduing an armed attacker (Briggs & Kennedy, 2016). It determined that compared to a no-fight scenario having unarmed individuals fight the attacker reduces casualties significantly (Briggs & Kennedy, 2016).

Finally, one unpublished work using ABMS conducted exploratory research into responding officer formations (Kristopher D. Davis et al., n.d.). It looked into the effectiveness of different formations like "T" and "Diamond" that officers can use when closing with a shooter and varying team sizes from single officers to four-man teams. The main conclusion of this paper was that a formation-based four-man responder team was the most effective at eliminating the threat with the lowest casualties to officers. Of course, with all these ASI models, certain factors are often not included, such as weapon lethality, shooter accuracy, and various behaviors of civilians and responders. This level of fidelity might help in creating more accurate models, allowing for better policy adoption and better decisions to be made in the future. As mentioned above, these models need to be considered within their intended context. They should support policy making, not mandating in the absolute what policy to adapt.

Overall, much has been done with ABMS in the field of ASI research. The author predicts that we will continue to see impactful work in the field given rapidly expanding machine learning (ML) technologies and the constant improvement of supporting tools and technology. As ABMS becomes more mainstream in academia and industry, more decision makers will be exposed to its strengths and weaknesses. They can therefore use it to make a better-informed policy decision.

2.6 Machine Learning

Within Machine Learning (ML), there are several recognized methods that are used to solve a specific problem set. Supervised learning, which can be used to predict values or classify objects by using labeled data to train the model, is the first and most commonly used method. A model that predicts the price of a car or classifies an email as spam is a good example (Géron, 2019). Unsupervised learning can broadly be thought of as pattern recognition, detecting clusters of objects in data without having the benefit of labels on each observation. Examples include visualization algorithms (Géron, 2019). Semi-supervised learning is, as the name suggests, a combination of both supervised and unsupervised algorithms (Géron, 2019). The model can be trained on labeled data yet is also capable of detecting patterns that can then be labeled based on further input. An example is a photo-hosting service that can detect different people in photos, yet might not have the labels for each detected person in a picture (Géron, 2019). The above algorithms require the developer to utilize many different approaches to train the algorithms to produce a useful model.

Lastly, we come to reinforcement learning, which is quite different from the methods above. Reinforcement learning (RL) uses an agent as the learning system that can observe and is influenced by the environment and can take actions to receive rewards in return for those actions (Géron, 2019). As Géron (2019, p. 14) states, “It must learn by itself what is the best strategy, called a policy, to get the most reward over time. A policy defines what action the agent should choose when it is in a given situation.”. This form of ML could be ideal for agent-based modeling and simulation (ABMS), and in particular for active shooting incident (ASI) researchers. Having a learning agent integrated into a simulation, in some way, might yield important information.

Within the last few years, RL-based software tools and models have seen enormous growth while achieving great feats, not just for the scientific community. The Company DeepMind, now owned by Google, created AlphaGo, a neural network-based RL program that defeated Lee Sedol, the world’s best Go player, in 2016 (Somers, 2018). Then, the company created a more powerful Artificial Intelligence (AI) AlphaZero (formally AlphaGo Zero) that is capable of learning to win numerous two-player games such as Chess and Go without any knowledge of the games, other than its rules (Somers, 2018). AlphaZero is able to beat all previous AlphaGo versions by being its own teacher through playing against itself, free of other inputs (Silver & Hassabis, 2017). Adding to the success, the AlphaStar AI was able to decisively defeat a professional StarCraft gamer in a 5-0 competition held under professional match conditions (The AlphaStar team, 2019). The game of StarCraft is a highly complex real-time strategy game that requires planning, situational awareness, strategic thinking, and constant tactical actions in order to win.

The fact that a machine could defeat a human player in a real-time environment opens up the possibility of integrating this type of AI into numerous other real-time models and applications. Also notable is OpenAI, an artificial intelligence company that seeks to advance digital intelligence (OpenAI, n.d.-a). OpenAI has created many useful tools and libraries, such as OpenAI Gym, a tool which allows for the training and comparing of RL algorithms (OpenAI, n.d.-b). The organization has been active in advocating for safe AI application to the benefit of all (OpenAI, n.d.-a). Its most notable accomplishment is creating OpenAI Five that, similar to AlphaStar, was able to learn to play an e-sports game and defeat the world champions as well as win 99.4% of over 7000 games played against professional gamers all across the globe (Berner et al., 2019).

This marks a major advance in RL and AI as a whole, given the complexities of the game of Dota2. According to OpenAI Berner et.al (2019, p. 1), the game of Dota2 involves two teams of five players facing off on a square map and "presents challenges for reinforcement learning due to long time horizons, partial observability, and high dimensionality of observation and action spaces". One can see that this problem is considerably more complex than the StarCraft game discussed above, given the number of independent players involved.

The main point to note is that RL-based AI is now capable of operating successfully in extraordinarily complex multi-agent environments and achieving victory against human opponents in real time. This level of sophistication might help researchers solve problems by integrating RL with ABMS for ASI research.

2.6.1 Reinforcement Learning

Sutton and Barto (2018, p. 1) state that "Reinforcement learning is learning what to do—how to map situations to actions—so as to maximize a numerical reward signal. The learner is not told which actions to take, but instead must discover which actions yield the most reward by trying them.”. Reinforcement learning (RL), and its related deep reinforcement learning (DRL), is an approach more focused on goal-oriented learning from interactions than other forms of Machine Learning (ML) (Sutton & Barto, 2018). As Géron (2019, p. 14) explains, the machine “… must learn by itself what is the best strategy, called a policy, to get the most reward over time. A policy defines what action the agent should choose when it is in a given situation.”. When the term “deep” is preceded by any form of ML, it is implied that a Deep Neural Network (DNN) architecture is used to train the model to achieve its purpose.

The problem of RL is formalized through an incompletely known, or partially known, Markov Decision Process (MDP) (Sutton & Barto, 2018). The MDP is intended to include just three aspects, the sensation (observation), action, and goal, and is hence fundamental to reinforcement learning (Sutton & Barto, 2018). RL is considerably different from supervised and unsupervised learning, two well-studied and used fields in ML. Both supervised and unsupervised learning do not require an environment to train on, like RL, and these techniques do not seek to achieve a goal per se. One could say about RL, as Sutton and Barto (2018, p. 2) do, that “In uncharted territory—where one would expect learning to be most beneficial—an agent must be able to learn from its own experience.”. This makes RL notably different, and in many ways more dynamic, in the field of ML. It is, therefore, often used to solve a different kind of problem set than supervised or unsupervised learning. One of these problem sets is the main focus of this dissertation.

2.6.2 Deep Reinforcement Learning and Agent-Based Modeling and Simulation

Agents can be independent entities within the simulation, interacting with their environments and each other. Further, deep reinforcement learning (DRL) is at its foundation the process of having an agent learn by observing the environment, taking an action based on the state of the environment, and then receiving a reward based on the quality of the action. The ultimate goal of the agent is to maximize its reward. This training loop leads to a policy that the agent can use to achieve the goal set out for it in the simulation environment. This process lends itself perfectly to agent-based modeling and simulation (ABMS), and the majority of DRL research is conducted within an agent-based simulation environment. One paper used agent-based modeling and DRL to discover the proper policy for optimizing the dosing of medicine for cancer patients (Jalalimanesh et al., 2017). The researchers used the NetLogo software and Q-learning to discover the optimized policy. NetLogo is one of several agent-based modeling tools available to researchers and industry (Wilensky, 1999).

Another application of reinforcement learning is in traffic control. One paper used Q-learning to discover a policy that outperforms traditional models, such as the Webster-based pre-timed signal control (El-Tantawy & Abdulhai, 2010). Further use of ABMS and DRL can be seen in stock market investments. One study developed an agent-based model using DRL to mimic professional trading strategies (Chen et al., 2018). Chen et.al (2018, p. 1) state that “The concept of continuous Markov decision process (MDP) in RL is similar to the trading decision making in financial time series data.”. The paper concludes that a reinforcement learning agent can successfully imitate the expert’s trading strategies at around 80% of the time (Chen et al., 2018). Another application of ABMS and DRL is in the energy market. One study developed a multi-agent-based system to balance supply and demand for electricity and discover the ideal prices, depending on the environment (Kiran & Chandrakala, 2020). It is interesting to observe that in a complex environment such as the energy market, reinforcement learning agents are able to look a day ahead in the market or in real-time and achieve high net earnings (Kiran & Chandrakala, 2020).

Outside of the sciences, ABMS and DRL are being used to accomplish great deeds. In the field of digital entertainment, video games are taking advantage of DRL to enhance their products like never before. For example, Unity, one of the premier game development platforms, has released ml-agents, a tool which allows developers to use reinforcement learning to train agents to act autonomously within the gaming environment (Unity, n.d.-a). Though their product is focused mainly on developers of video games, there is a viable application for researchers in other fields. This tool allows for the interfacing of Unity and a Python environment, enabling the creation and running of DRL networks through machine learning libraries such as Google's TensorFlow (Paris Buttfield-Addison, 2019). This combination of tools furthers the potential research that can be conducted in any field seeking to explore the power of ABMS and DRL. The Unity ml-agent toolset features a robust array of DRL algorithms to use and hyper-parameters to tune, even without the use of one's own custom neural network (Ervin Teng, 2019). As one can imagine, this tool will allow an easy and efficient way to build environments to train DRL agents for any research field. Further, Unity provides an ml-agents wrapper for the popular reinforcement learning interface OpenAI Gym (Unity, n.d.-b). This is significant given the popularity of OpenAI Gym with reinforcement learning researchers and hobbyists alike. Some examples of Unity ml-agent use are training a car to drive around a racetrack without hitting a wall or training a robot to push a box into a specific area (Paris Buttfield-Addison, 2019). Of course, there are many more complex applications of this tool, given the nature of ABMS with DRL.

One of the most powerful AMBS modeling tools is the AnyLogic software system (AnyLogic, n.d.). This system allows developers to create models from low to high fidelity and produce data for analysis. A complementary tool for use with AnyLogic, for the sake of combining AMBS with DRL, is Pathmind, a software as a service (SaaS) tool (Pathmind Inc., 2021) (Farhan et al., 2020). Together with the simplicity and power of AnyLogic, Pathmind allows researchers to easily create observations, actions, and the reward function, all through a simple interface. Then, one can train the policy within the Pathmind system by providing the AnyLogic model and the various parameters needed. Monitoring the training process is also made simple with a web interface so the researcher can ensure that training is progressing as expected. Once the policy is trained, it can be integrated into the existing model with a few mouse clicks and a simulation run to gather data.

The author's intent at focusing on these tools is to demonstrate how robust and viable DRL and ABMS tools are and that these technologies are ready to be applied to other research fields. Quite a lot has been done with ABMS and DRL in business and research in the last few years. This is partly due to the increase in available processing power, particularly for individuals and small research teams, and the improvement and release of many new tools and technologies related to DRL and machine learning in general.

2.7 The State of ASI Research Using ABMS and DRL

The author has found precious little research addressing active shooting incidents (ASIs) combining agent-based modeling and simulation (ABMS) and deep reinforcement learning (DRL). One study uses DRL in a simulated environment to control Unmanned Aerial Vehicles (UAV's), commonly referred to as "drones", to take a picture of a shooter's face for identification (Tzimas et al., 2020). Though definitely useful, this is not the type of research the author has so far discussed regarding active shooting incidents, which is somewhat "force on force", where agents act in adversarial capacity. In line with this, one paper describes the use of reinforcement learning to develop tactics in air combat (Piao et al., 2020). Though not directly related to ASI research, it does have an element of discovering policies for agents to attack and defend against each other. This lends credence to the legitimacy of combining ABMS and DRL for ASI research.

2.8 Summary

Active shooting incidents (ASIs) have steadily increased in volume and lethality over the last few decades. There have been many instances of lethal attacks against schools, churches, businesses, and other public places. Most attacks occur in static indoor spaces, yet open-air venues and open spaces are particularly vulnerable. There have been several highly lethal and destructive attacks on open-air venues in the recent past. Therefore, these vulnerable spaces need more attention from law enforcement and researchers to better secure them in the future. Much has been done in the last two decades to help mitigate the damage done by active shooters. Police tactics have improved since the Columbine High School shooting, and civilian actions have been derived that give potential victims a higher chance of survival during an incident.

A steady output of quality research in the field covering active shootings continues, and a lot of it uses computerized modeling to assist in producing data for analysis. Agent-based modeling and simulation (ABMS) has been used in the recent past to help discover policies for decision makers to adapt in defense against active shooters. Many tools such as AnyLogic are available for researchers to use in their quest for a better understanding of the problem. With the maturity of machine learning (ML) comes an opportunity to use this technology and the associated tools to push the envelope of ASI research. Using powerful products like Pathmind, which uses deep reinforcement learning (DRL), together with AnyLogic, could yield advances in the way we think about defending innocents against violent attackers.

CHAPTER 3. RESEARCH METHODOLOGY

3.1 Overview

The research methodology chosen for this study supports gaining insight to answer the stated research questions. Agent-based modeling and simulation (ABMS) was chosen to produce data for analysis. ABMS has proven to be valuable in the field of active shooting incident (ASI) research, and the author's familiarity with the technologies and concepts is of value to the project. In the analysis of past events, ABMS has yielded policies that can be adapted to mitigate active shooter impact. The data produced by a properly validated ABMS model is as close to human behavior and as realistic as possible without putting actual people at risk. Work can also be conducted at a low cost to the researcher.

Together with new technologies like reinforcement learning (RL), ABMS offers the possibility of cutting-edge experimental research into the field. This has the potential to discover new and effective policies that decision makers might adapt to protect life and property. Further, new technologies can be examined in a safe environment to deduce their value without having to build prototypes, further reducing costs and assisting leaders in making decisions. The researcher chose a specific set of technologies to accomplish his task, and a specific scenario to model that would effectively answer the research questions.

3.2 AnyLogic Modeling Software

AnyLogic is a Java-based modeling software tool that allows modelers to visually construct their work by connecting logical components in a graphical user interface (GUI). It is used by industries such as manufacturing, mining, transportation, rail logistics, and many others (AnyLogic, n.d.). It is also prevalent in academia and research organizations. The software system is flexible enough for individuals to use system dynamics, discreet and agent-based modeling, or a combination of all three. Its ability to allow a researcher to easily change parameters in the model and run the same scenario with these new values is a key benefit, enabling easy and flexible experimentation to produce useful data for analysis. Further, the ability to create 2D and 3D environments with ease and accuracy, mimicking real-world places and spaces, is of enormous value.

The software is backed by many large and well-known corporations and institutions and is bolstered by many white papers and case studies. The product is trusted by the author and his colleagues due to repeated use in research. This work has listed many active shooting incident (ASI) models created with AnyLogic, yet there are other fields that use the software for research. For example, one team modeled a regional hub reception center for evacuation after a disaster (A. Kirby et al., 2012). Another set of researchers used AnyLogic to model the interception of an unmanned aerial vehicle (UAV) (Cline & Dietz, 2020). The AnyLogic website demonstrates through various publications that the underlying scientific software is accurate and reliable (AnyLogic, n.d.). Users can be confident that the software produces accurate calculations at runtime. These benefits are the reason the AnyLogic modeling tool was chosen for this project.

3.3 Pathmind Software

Together with AnyLogic, the author used Pathmind, a deep reinforcement learning (DRL) system that integrates with AnyLogic. As discussed above, reinforcement learning (RL) is a powerful methodology to solve a set of problems. Pathmind provides the ability to use this power to experiment within an agent-based model to help mitigate the active shooter problem. The principal reason to use Pathmind for the DRL solution is its ease of use and ability to integrate seamlessly into an established agent-based modeling tool. It is simple to designate observations, create the reward function and receive actions from the trained policy using Pathmind's interface. It consists of easily understood and implemented code tie-ins within AnyLogic via a plugin. A modeler simply has to make the "connection" with the AnyLogic model, specify the observations and the reward parameters, then update a few settings in the Pathmind helper plugin. After this, the model is uploaded via a web interface to train the policy. An example of using the web interface to set up an experiment to train a policy with Pathmind is visible in figure 3.1 below. There is no need to develop a custom neural network with various amounts of nodes and hidden layers or worry about the many algorithms and hyperparameters or their tuning. Of course, the downside of this simplicity is the loss of detail maintained with the neural network. However, the drawback has no significant impact considering the breadth of DRL work and researcher experience with the technology. More than that, it is of great value to spend the majority of time developing the model and running experiments instead of debugging and fiddling with the neural network.

The author can be confident in the capability and power in the Pathmind technology, given what is being used "under the hood". Though the exact size and makeup of the neural network is proprietary, Pathmind has published various details that provide insight as to the fidelity of the system. For one, Pathmind uses the proximal policy optimization (PPO) algorithm to optimize the hyperparameters and population-based training (PBT) to train the model (Farhan et al., 2020). Both of these algorithms and methods are well-established and well-known. They are both effective and used in many ML and DRL systems to train a model and fine-tune the hyperparameters. Pathmind abstracts the details of training and fine-tuning to the benefit of the user. As Farhan et al. (2020, p. 8) state, "Pathmind automates this process to simplify a user's experience on running RL experiments.".

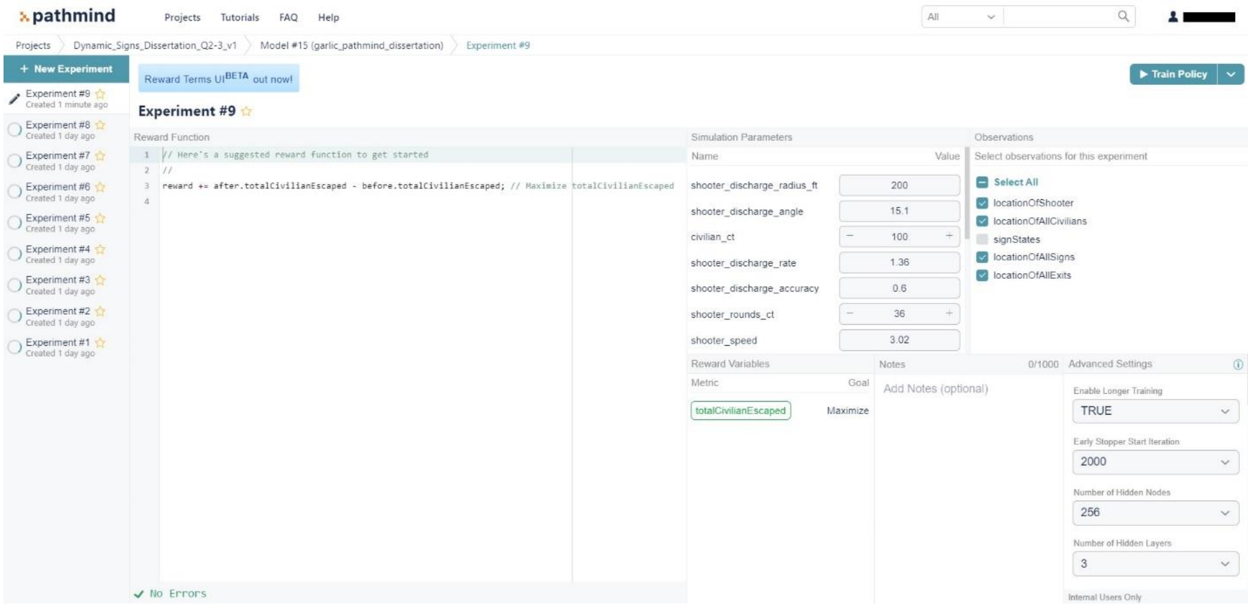

Figure 3.1 - Example Pathmind Experiment Setup

3.4 Gilroy Garlic Festival Model

The author discussed the details of the Gilroy Garlic Festival incident in Chapter 2 and it is an integral part of this work. However, the reader should be aware that the research conducted herein is not a case study of the shooting. It is simply used as a basis to conduct research and answer the specified questions. For further detail on the shooting itself, Frantz (2021, pp. 58-59) has detailed the background of the annual festival as well as the actual incident.

3.4.1 AnyLogic Model Foundation

The model design and implementation are based on the work performed by a researcher who investigated the Garlic Festival active shooting incident (ASI) (Frantz, 2021). As stated by Frantz (2021, p. 60), “The model design was based upon a diagram of the Gilroy Garlic Festival at Christmas Hill Park in Gilroy, California.”. Frantz asked questions about the number of police officers that should be available to minimize casualties and what positions they should take within the venue to maximize patron survival (Frantz, 2021). Further, he investigated how many exits were most advantageous and where to place them. Due to these efforts, a full, feature-rich, and validated model exists on which to base new work. Though extensive work needed to be conducted to modify agent logic and integrate Pathmind deep reinforcement learning (DRL), the time saved by having the basic layout of the festival grounds as well as the validation parameters allowed the author to perform his work without spending time on the basic structure. Further consideration of the reference is recommended to discover the details of the foundational design the previous researcher implemented (Frantz, 2021).

3.4.2 Foundational Model Design Details

Some of the details of the foundational model design are worthy of note. The previous researcher based his environment on diagrams of the Garlic Festival open-air venue and produced a representation of it within AnyLogic.

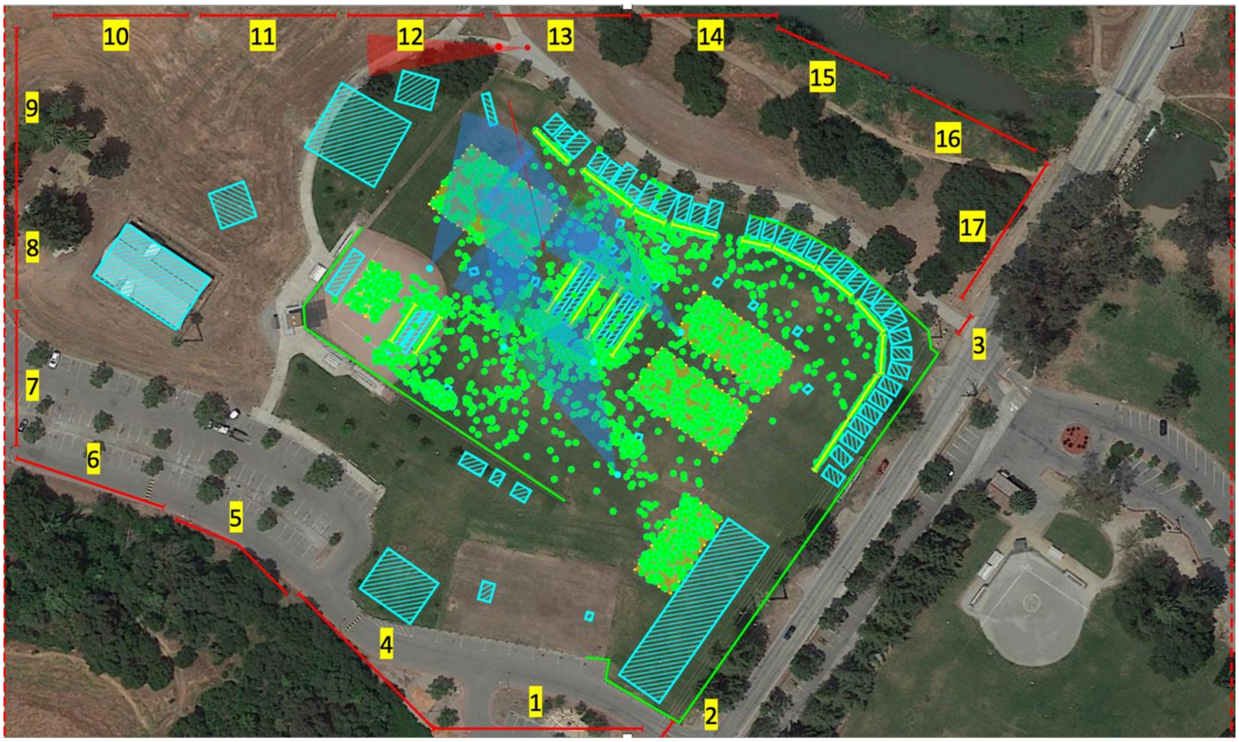

Figure 3.2 - Model with 17 Exits
(Frantz, 2021)

Represented in this image is the real-world terrain, gathered from a mapping service, and the all-important overlays within AnyLogic that represent the vendors (light blue rectangles), tents (yellow rectangles), benches (yellow lines), and other objects that were present at the time of the shooting. Also represented are exits (red lines) and fencing (green lines) surrounding the main festival grounds. This screen capture demonstrates a model run with thousands of civilians, the shooter, and the police responders. Though discussing the actual logic of the base model would be useful in any other effort, it is not important at this stage due to the fact that the logic for the new model has changed so significantly it would not be useful to highlight the foundational logic. Also, certain components like police responders were removed since they were not necessary for this research. Only the environment stayed largely consistent as well as the shooter logic, which will be discussed below in the context of the new and improved model.

3.4.3 Foundational Model Parameters

The parameters used by the model are present to ensure appropriate data collection. The foundational model uses the parameters listed below in Table 3.1. They were created by the previous researcher and used to answer his research questions and therefore are of value to this follow-on work (Frantz, 2021). All but the police responder-related parameters are used in the new model. The exact parameters used from the table below, as well as new ones added for data collection in the new model, will be listed in a separate section.

Table 3.1 - Model Parameter Description
(Frantz, 2021)

Output Name	Data Type	Descriptor
model_runtime	integer	Model duration in seconds
shooter_discharge_radius	double	Shooter weapon discharge radius in degrees to engage civilians and police
shooter_discharge_angle	double	Shooter field of fire in degrees
shooter_discharge_accuracy	double	Percentage of shooter rounds fired that strike target
shooter_discharge_rate	double	Discharge interval of shooter
shooter_target_ct	integer	Total number of civilians and police
shooter_casualty_ct	integer	Total casualties caused by shooter
shooter_speed	double	Shooter movement in ft/second
civilian_ct	integer	Total number of civilian agents within the model
police_discharge_radius	double	Police weapon discharge radius in degrees to engage shooter
police_discharge_angle	double	Police field of fire in degrees
police_discharge_rate	double	Discharge interval of police
police_discharge_accuarcy	double	Percentage of police rounds fired that strike the shooter
shooter_rounds_ct	integer	Total rounds fired by shooter
shooter_duration_end	long	Model ends when shooter is neutralized by police

3.4.4 Foundational Model Validation

The model had been validated by the previous researcher. It is important to point out some specifics of this validation to confer some confidence as to the validity of the foundational model used for further work. The researcher validated the model by deriving the proper parameter values listed in Table 3.2 below. All values were used except those related to police responders and the number of civilians present at the time of the shooting. This difference will be discussed in detail when the author outlines additions and changes to the new model. The process used is the same as other ASI researchers have used in other insightful models (Lee, 2019). It involves running the designed model with Monte Carlo runs until the various derived parameters satisfy the real-world results of specific variables, such as the duration of the shooting, the number of rounds fired, or the speed of the shooter and other persons present. These parameter values are listed as derived from the "Model". Historical parameters were used by the foundational model designer and are listed as "Historical". Further, values that were gathered from literature are listed as "Research". The process of how the validated parameters were derived is detailed by Frantz (2021, pp. 68-87).

Table 3.2 - Model Validation Parameters
(Frantz, 2021)

Parameter Name	**Type**	**Value**	**Source**
Shooter_speed	Fixed	3.02 ft/sec	Model
Police_speed	Fixed	3.1 ft/sec	Model
Civilian_speed	Fixed	3.28 ft/sec	Research
Shooter_discharge_radius_ft	Fixed	200 ft	Model
Shooter_discharge_angle	Fixed	15.1°	Model
Police_discharge_radius_ft	Fixed	200 ft	Model
Police_discharge_angle	Fixed	2°	Model
Police_historical_ct	Fixed	3	Historical
Civilian_ct	Fixed	3,290	Historical
Shooter_ct	Fixed	1	Historical
Shooter_discharge_rate	Fixed	1.36 sec/round	Historical
Police_discharge_rate	Fixed	3.33 sec/round	Historical
Shooter_accuracy	Fixed	0.6 (60%)	Model
Police_accuracy	Fixed	0.15 (15%)	Model
Shooter_discharge_ct	Fixed	36	Historical
Police_discharge_ct	Fixed	18	Historical

3.5 Expanded Model Design

This section discusses the necessary additions and changes made to the foundational model to produce data and train the deep reinforcement learning (DRL) policy. The expanded model is not intended as a case study of the Garlic Festival active shooting incident (ASI), as was the foundational model it was built upon. The expanded model is meant as a testbed to investigate and answer the research questions through experimentation.

3.5.1 Adjusted Model Parameters

The parameters used by the expanded model are listed in Table 3.3 below. They are a reduced set of parameters due to the limitations and delimitations mentioned in this paper. Notable changes are the civilian count reduced to 100 civilians, the unlimited ammunition count, and the 100% accuracy of the shooter. All three were changed so as to produce more meaningful results in the final model and to assist in training the DRL policy. Just as in the foundational model, the source of the values is from model validation, historical source, or literature research. One exception is the count of 100 civilians, which was chosen to reduce the training time of the DRL policy and to reduce the processor load on the model during Monte Carlo runs. The researcher does not believe that using the foundational model's value of 3290 civilians would contribute anything meaningful in experimentation but rather bog down the work.

Table 3.3 - Adjusted Model Parameters

Parameter Name	Type	Value	Source
Shooter_speed	Fixed	3.02 ft/sec	Model
Civilian_speed	Fixed	3.28 ft/sec	Research
Shooter_discharge_radius_ft	Fixed	200 ft	Model
Shooter_discharge_angle	Fixed	15.1°	Model
Civilian_ct	Fixed	100	Research
Shooter_ct	Fixed	1	Historical
Shooter_discharge_rate	Fixed	1.36 sec/round	Historical
Shooter_accuracy	Fixed	1.0 (100%)	Model
Shooter_discharge_ct	Fixed	unlimited	Model

3.5.2 Model Environment

The model environment is the physical space of the historical event that was recreated by the previous researcher (Frantz, 2021). Landscape, buildings, and structures, as well as the scale of such objects, are all worthy of mention. Very little needed to be modified from the foundational model. All of the buildings and structures were maintained at the exact location, size, and scale. The only significant change is the modeling of exits. Each exit in the foundational model needed to be split into three separate exits. This was necessary given how the civilians detect dynamic signs and the exits (discussed in detail below). For the civilians to have a better chance of detecting the exits, they needed to be split apart. This is because of a technical issue within the AnyLogic software and how it places and keeps track of TargetLine classes. Each line has a center point that it uses for the X/Y coordinate in the plane. This coordinate is used by the civilian detection algorithm to “see” the exit line. If the detection algorithm that uses a field of view mechanism does not cover the center of the TargetLine, the exit will not be detected. Therefore, the fidelity of the model was increased so as to better model the civilians being able to detect exits. The new model environment is displayed in Figure 3.3 below. Also of note are the black Attractor-class objects present in the model. These are the representation of the dynamic signs that are crucial to answering the research questions. These signs and their modeling will be discussed in detail in the following section.

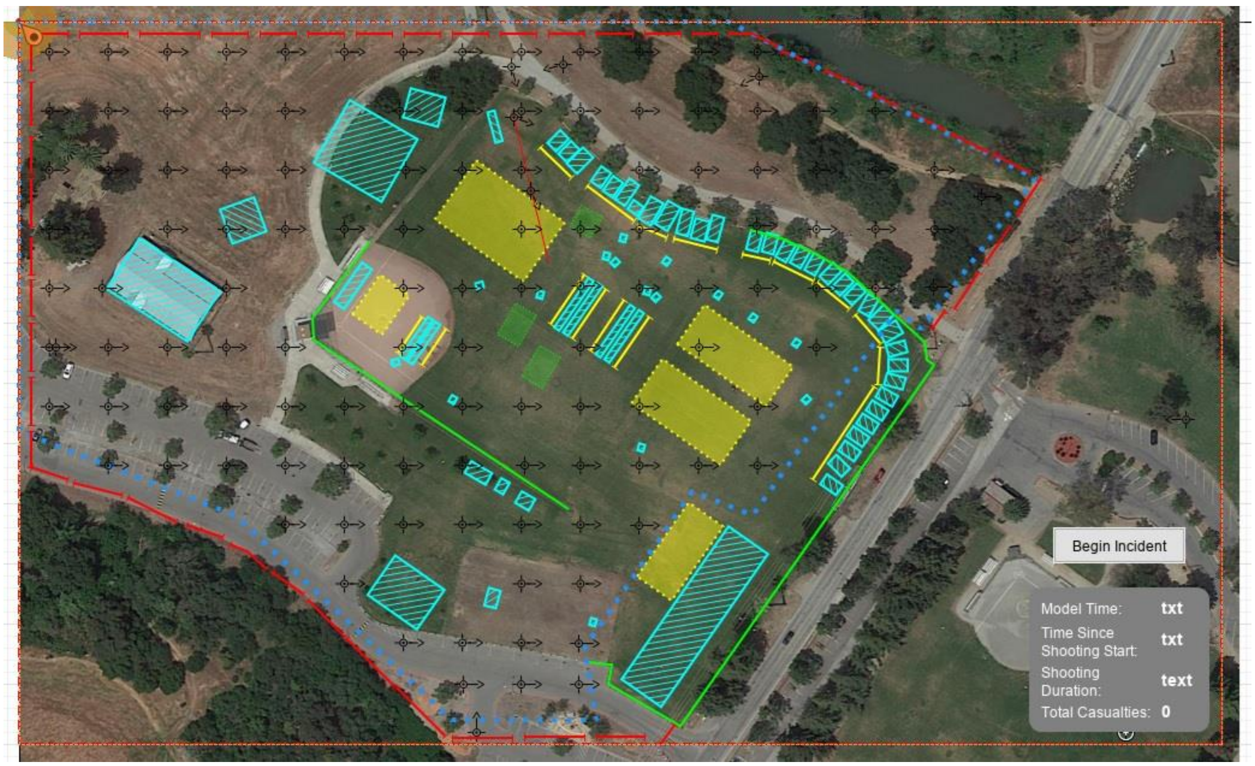

Figure 3.3 - Adjusted Model Environment

3.5.3 Added Model Parameters

New parameters were required to answer the research questions. These parameters are mainly associated with the civilian agents in the model and are critical to the successful training of a DRL policy. One parameter was inserted to support the spawning of the shooter agent in one of the cardinal directions. These are north, south, east, and west. It was an important property to add to obtain knowledge on the DRL policy's ability to train with this feature included. Though the author believes he could answer the research questions with a single spawn location, being able to change the spawn location and train a policy for each adds more value to the research and more confidence in the results. The new parameters are listed below in Table 3.4.

Table 3.4 - Added Model Parameters

Parameter Name	**Type**	**Value**	**Source**
civilian_view_angle	Fixed	360 °	Model
civilian_view_angle_run	Fixed	120 °	Model
civilian_view_radius_ft	Fixed	120 ft	Research
shooterSpawnLocation	Variable	Cardinal Direction	Model

3.5.4 Dynamic Signs

The core of the expanded model is the dynamic signs added to the environment, which are intended to guide civilians to a safe exit, away from the shooter. The signs are a real-world product under development that can be deployed by a venue in advantageous positions around an open space. They will include various features to gain the attention of civilians in an emergency situation and direct them to a safe location. One of the intents of the author's work is to evaluate the efficacy of a DRL policy to control these signs during an evacuation. Only certain signs should activate based on the location of civilians and the location of the threat, therefore guiding a maximum of civilians away from the threat. Frantz (2021, pp. 116-128) details the design and intended use of these signs.

To properly model the signs, the author used AnyLogic Attractor objects on the overlay. The distance between each sign is 60 feet, half the distance of the civilian's ability to see via its field of view arc. A matrix of these signs was automatically placed by the AnyLogic system with the given distances. A total of 119 signs were placed on the overlay after manually removing a few that were within modeled buildings and structures. These were removed, given that the AnyLogic system would throw errors if a civilian agent were not able to reach the modeled sign, which would occur if the attractors were placed within objects. A few signs were also removed from civilian wait areas, such as the modeled benches and tents, since this placement would not make much sense in a real-world situation. The overall intent of the placement of the signs was to inundate the area so that the Pathmind DRL policy could train and learn which signs were significant and which were not.

It was hoped that this produced an initial answer to the first research question, even if many more signs were present than realistically needed. If the trained policy improved survivability, it would be a matter of more effort to reduce signs in the environment while still maintaining increased survivability. This experimentation is discussed in detail in later chapters. The placement of the signs is shown in Figure 3.4. The signs, represented as AnyLogic Attractor objects, are highlighted in magenta on the figure. They are spread within the blue dashed line, which is an area within the venue, surrounded by the red exits. It is worth repeating that the civilians will spawn within the yellow areas and then randomly move toward another yellow area or a yellow line, representing the festival tents and benches. Once the shooter spawns, the incident begins, and the signs are intended to guide the civilians to safety, represented by a red exit line.

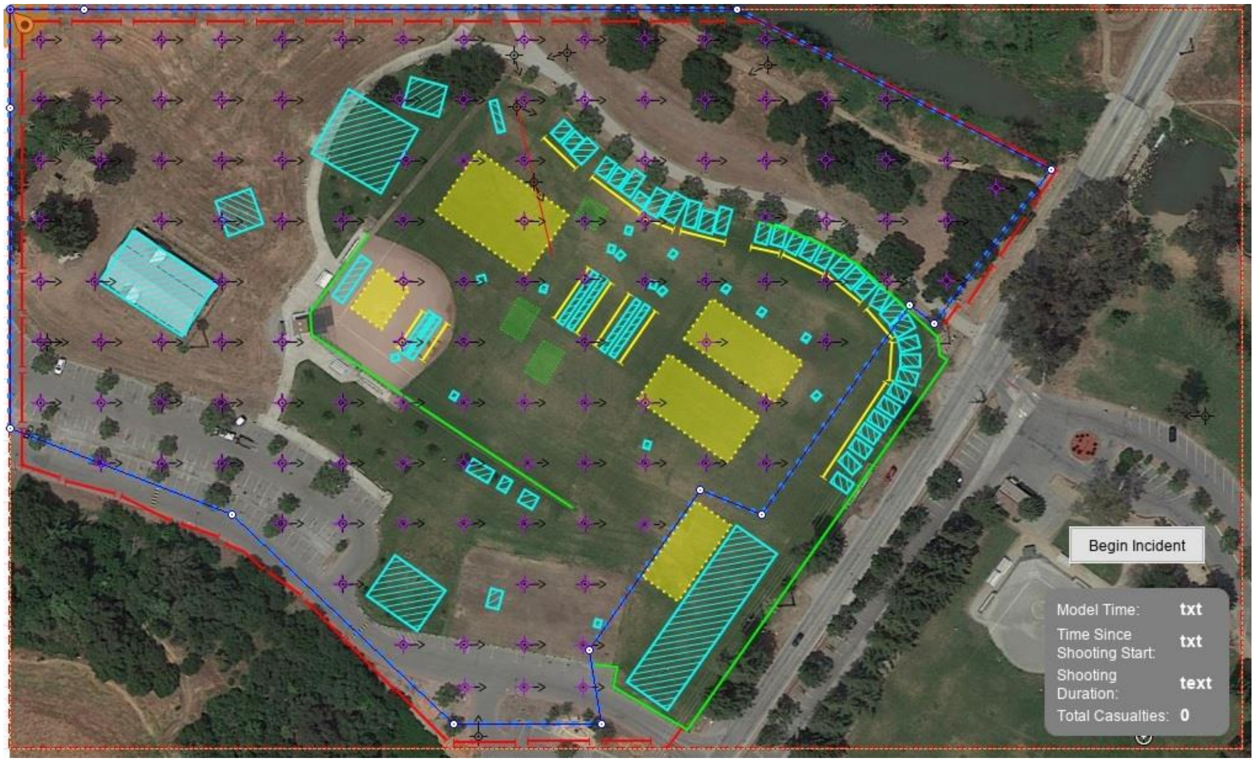

Figure 3.4 - Base Dynamic Sign Placement

During simulation runtime, the signs are represented by ovals that are colored red when they are inactive or green when they are active. This gives the user a better understanding of the state of a sign in real-time. Also of note is that the shooter is represented by a dark red circle and the civilian festival patrons by a light green circle. The shooter's field of view is also displayed. See Figure 3.5 below for a screenshot of the running simulation.

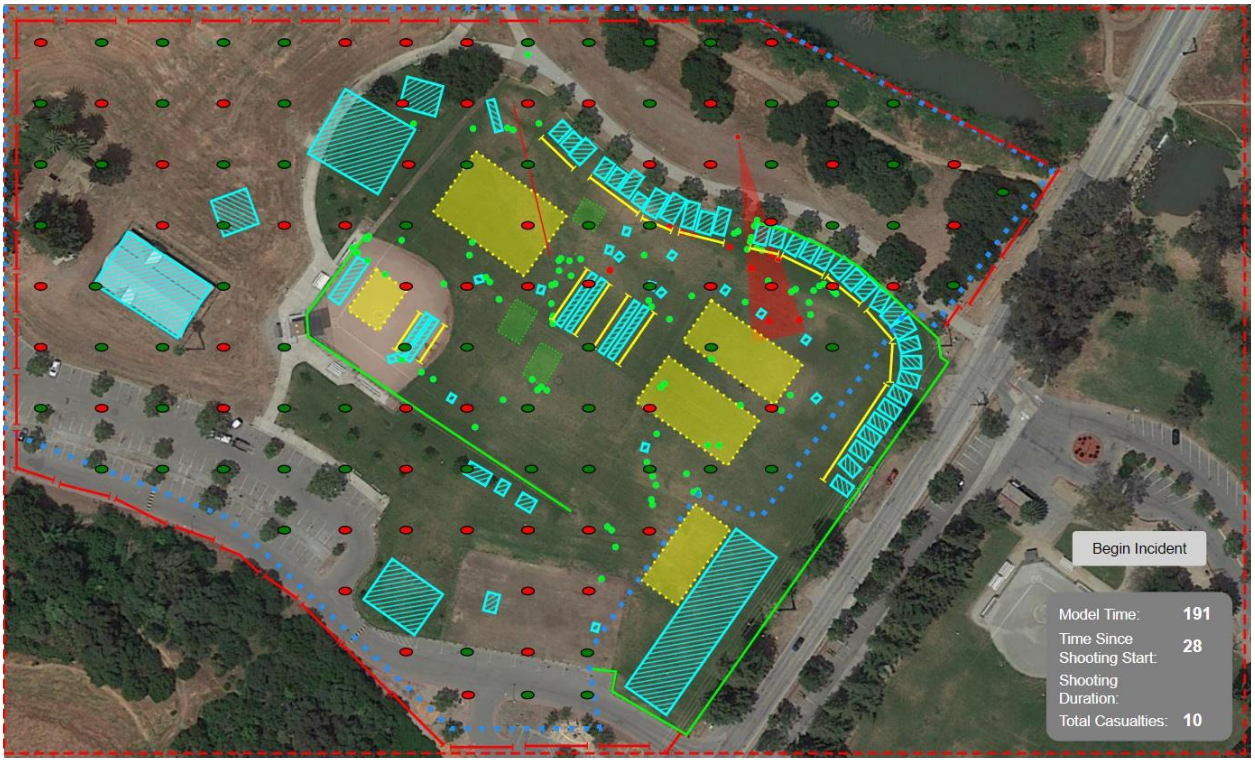

Figure 3.5 - Dynamic Signs During Simulation Run

3.5.5 Shooter Agent Design and Logic

The shooter design and its associated logic within the AnyLogic system required updates and improvements from the foundational model. Firstly, a minor fix to work around a small fault in the software was introduced, which has no bearing on the output data or the final results. This was a short delay of one millisecond for the shooter to reset its location in the underlying code. It was necessary to make the model stable during Monte Carlo runs. Secondly, the logic was changed to use different AnyLogic blocks to accomplish the same task. Instead of “Wait” blocks, the new logic used “GoTo” blocks to move the shooter towards the civilian targets. This use of the different blocks makes the code easier to understand. A simple check was added to assess if the shooter is still in the process of engaging. This part was a re-engineering of the logic to assess if the shooter was deceased or not. Though this was not needed since there are no responders in the scenario and the civilians do not fight, it would still be of benefit to maintain in the current model for the ease of future expansions. None of the additions and changes have any effect on the basic logic that was used to validate the shooter parameters in the foundational model. Even if so, it would have little bearing on this research since it is not necessarily reliant upon these parameters to produce useful and analyzable data. The schematic of the shooter logic is shown in Figure 3.6 below.

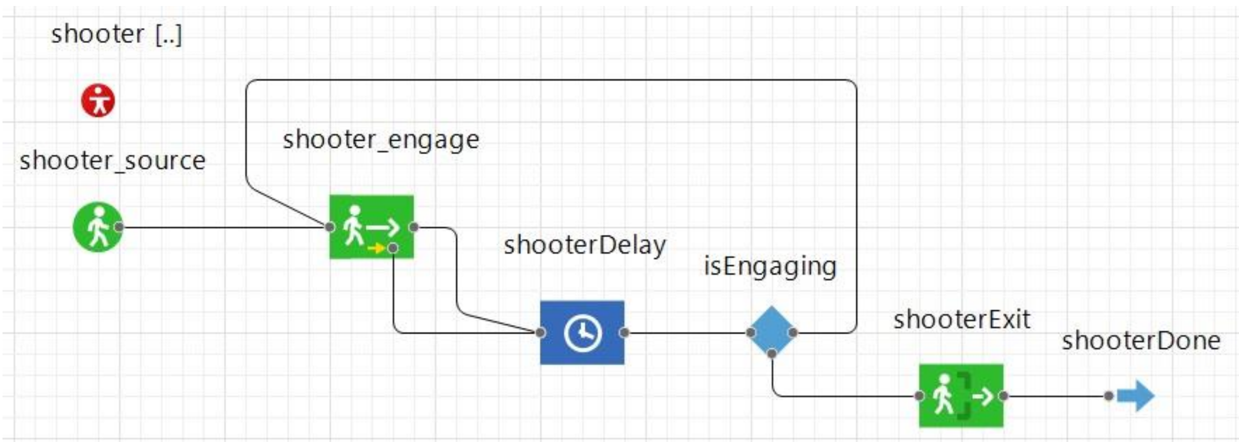

Figure 3.6 - Shooter Logic

The shooter type within AnyLogic did not require any significant changes, other than cosmetic adjustments and minor cleanup of the Java code. The shooter type schematic is displayed in Figure 3.7. For brevity, the schematic does not display the shooter field of view (FOV) arc that is used by the code to detect civilian agents within the system. However, this FOV arc is a central piece of the shooter agent's logic. It is used by the nearest_civilian_target() function (present on the Main canvas object) to detect any civilian agents within the FOV and then select one of those targets to engage. At this point, the shooter agents use the discharge() method to engage the civilian and assess if the civilian is a casualty or not. The most important change in this logic is that the shooter has a 100% accuracy in this model vs. the historical 60%. This was done to remove some fidelity in the model to make training the deep DRL policy more effective.

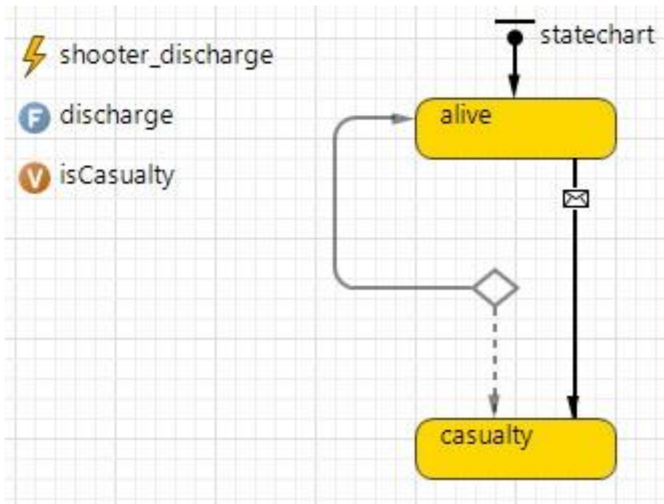

Figure 3.7 - Shooter Agent Schematic

3.5.6 Civilian Design and Logic

The civilian agent logic required extensive changes and additions to support this research, for the training of the Pathmind DRL policy in particular. To ensure that the civilian agents can detect the signs and exits in the model, a solution needed to be found to accomplish this in a way that provided enough fidelity yet was not so complex that it bogged down training and data collection. The solution was to use the same method that the shooter employed to detect civilian agents. An arc was added to the civilian agent that could be set to a certain radius and angle of view, creating a field of view (FOV) that properly models human vision. With this in place, code was developed to select the farthest sign within the agent's vision and move towards it while searching for any exits in view. Once an exit is found, the civilian will disregard any other signs and move to the safety of the exit. Of course, logic to assess if the civilian is a casualty was also refined in the new model. This logic is reflected in Figure 3.8. A "Delay" block is also present in the civilian logic, just like with the shooter agent. This is to overcome a minor issue within the system so that the model does not stall during Monte Carlo experiments.

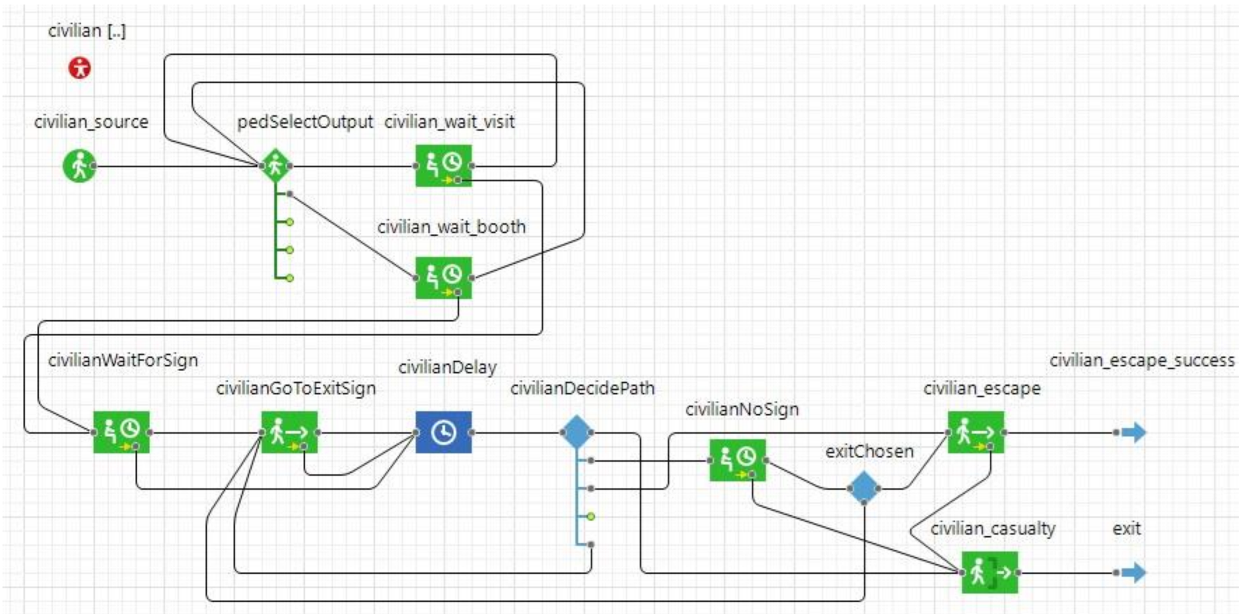

Figure 3.8 - Civilian Logic

Since the civilian logic is central to allowing the DRL policy to be properly trained, a simple workflow diagram is referenced to better demonstrate how the logic progresses, without needing to understand how AnyLogic works. Figure 3.9 shows the steps a civilian takes from the beginning of an incident until the civilian is either made a casualty or successfully escapes to an exit.

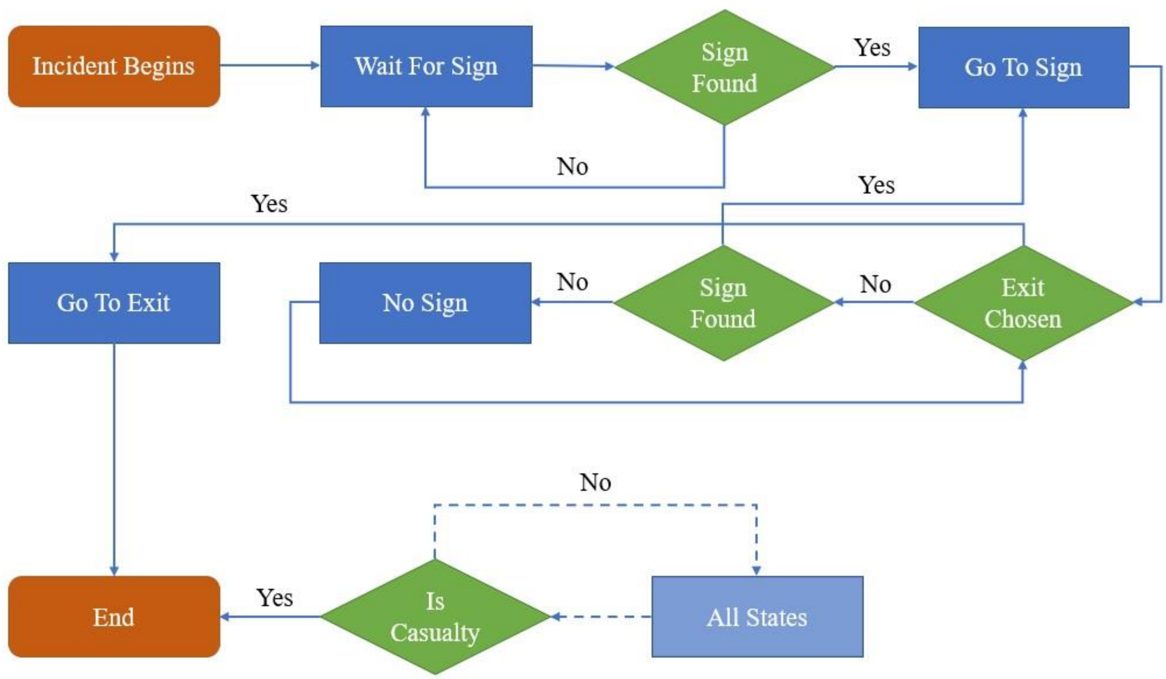

Figure 3.9 - Civilian Logic Workflow Diagram

The diagram above describes what the civilians will do after being randomly dispersed throughout the environment at specific locations, representing tents, benches, and vendors. This preamble to the incident is so that civilians are randomly dispersed throughout the simulated environment and is not relevant in the core logic needed to train the DRL policy. Once the incident begins, the civilian agent will wait for a sign to activate, to detect any signs within a 360-degree arc and a radius of 120 feet. When signs are found, the civilian will choose the farthest sign and go to it. Important to note is that the FOV of the civilian agent is now only 120 degrees with a radius of 120 feet. Once arrived within 30 feet of the sign, it will choose another sign and repeat the process. If a sign is not within the FOV of the civilian agent, once it arrives at the previous sign, it will wait at that location until an active sign comes in view. This is an important step since the DRL policy will have to recognize the state and activate a sign in the civilian's view. If during this process, any exit comes into the FOV of the agent, it will ignore any signs and run toward the exit and to safety. Once the civilian agent safely exits, it will be counted as a survivor. If it is engaged by the shooter and made a casualty in any of the states, it will exit out of the entire process and be counted as a casualty.

The civilian agent schematic below reflects the complexity of this agent type. Though the state chart is rather simple, many functions and class variables are needed to support agent operation and training within Pathmind. The author chose to encapsulate as much logic and data within the civilian agent type as possible. The code allows the civilian to adjust their field of view angle to provide 360-degree vision when the shooting incident begins, then narrow it down to 120 degrees. The 360-degree vision was implemented to model people's ability to look around and acquire their first sign. It was found that this helps better train the DRL policy. Once the first sign is chosen, the narrower field of view models humans moving forward towards the next sign. The 120-degree view angle was chosen to better train the DRL model and does not reflect any concrete human's peripheral or focused vision. The researcher will leave this level of fidelity and modeling to future work. The above logic is critical to allow the DRL to properly train and activate or deactivate certain signs in real-time. Detecting a sign within the field of view and walking towards it is the core of the logic that is integral to allowing the DRL agent to train and for the policy to guide civilians to safety. Other variables and functions were written to support this, as well as data collection. Figure 3.10 demonstrates the civilian agent type.

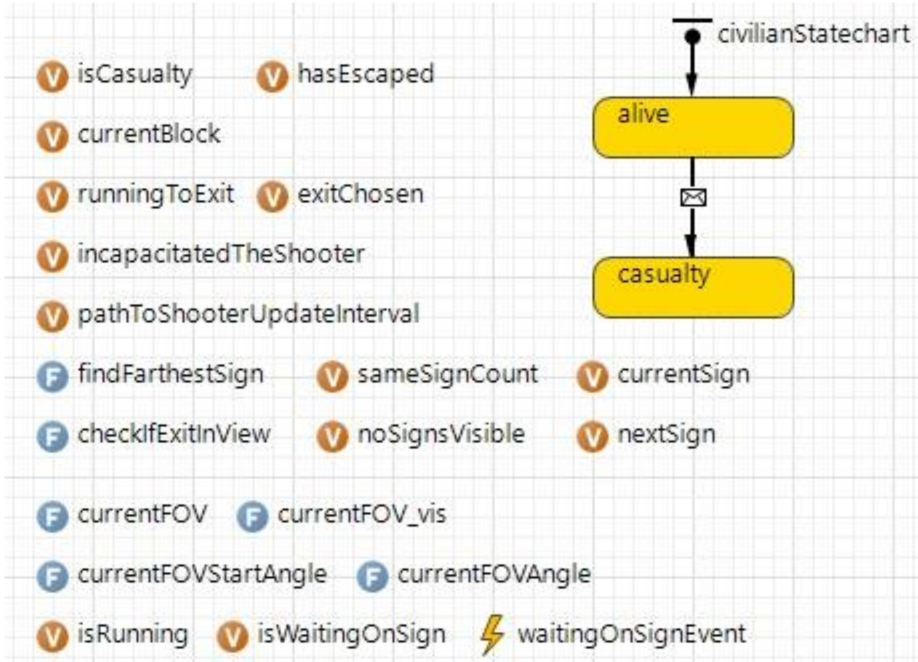

Figure 3.10 - Civilian Agent Schematic

3.5.7 Experiment Screen and Setup

To properly set up Monte Carlo experiments and efficiently produce data for analysis, the author created an experiment screen. This is a means to quickly set up the numerous experiments to produce comparable data. Further, it allowed the researcher to monitor the most important metrics during Monte Carlo runs in real-time, which proved valuable at detecting any errors early and saved time. The screen allows the user to set the various parameters needed for each run, such as where the shooter spawns or what kind of data needs to be collected. The author chose to collect two sets of data for comparison to the Pathmind DRL data that will be produced using a trained policy. Those sets of data were produced by two scenarios. One, where all signs are turned on, allowing the civilian agents to choose any one to attempt to escape, and two, where the civilian agents choose the closest exit to escape. These two scenarios give an understanding of what the outcome might be without any actively controlled signs in the environment. The experiment screen and what is displayed during a Monte Carlo run are shown in Figures 3.10 and 3.11 below, respectively.

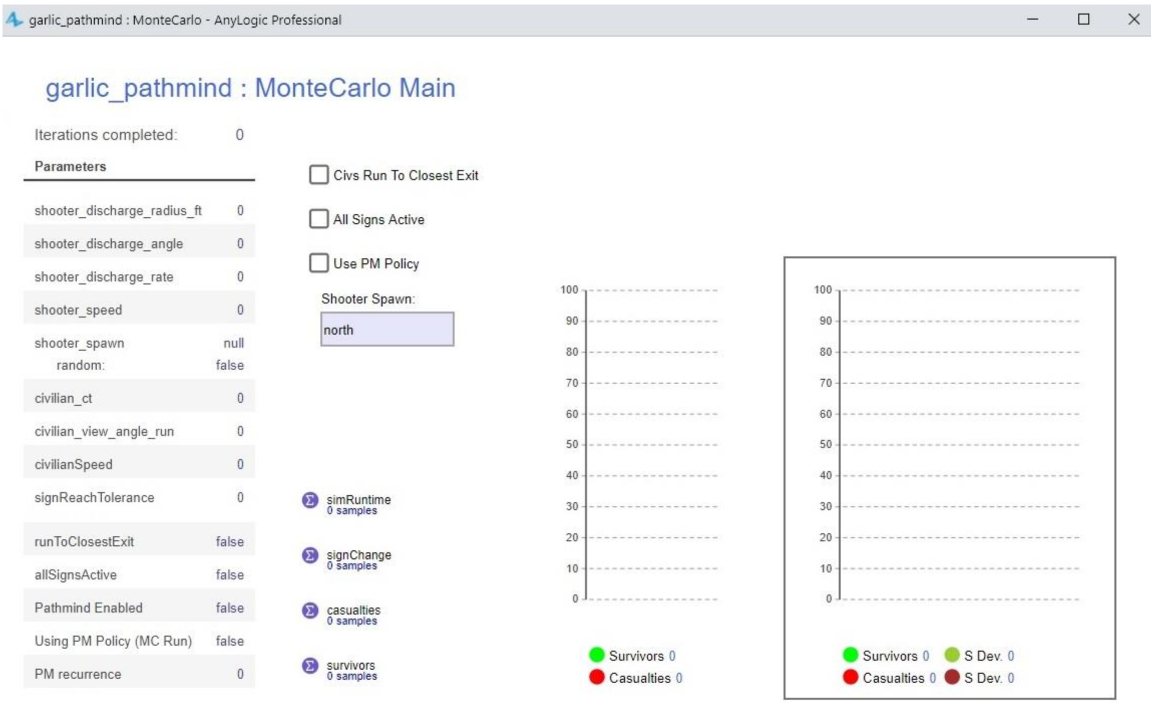

Figure 3.11 - Experiment Screen

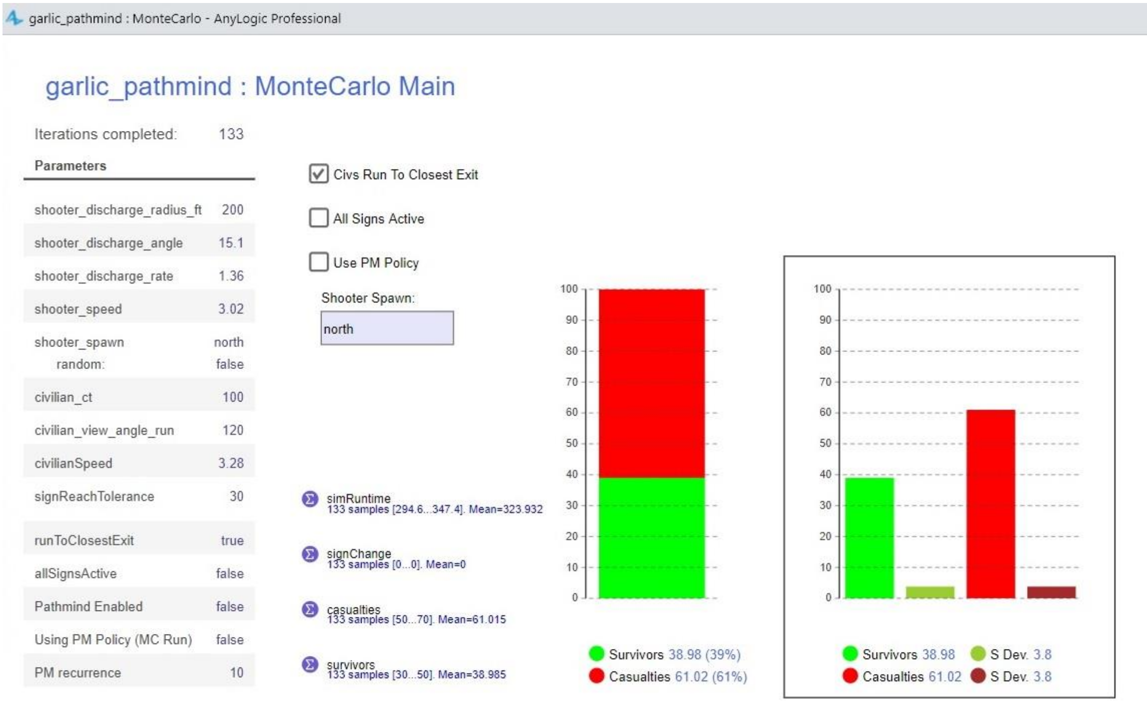

Figure 3.12 - Experiment Screen Running

3.6 Pathmind Integration and Use

The most important element of this experimental research is the deep reinforcement learning (DRL) capabilities of Pathmind within an AnyLogic-built environment. It is at the heart of this work, given the stated research questions. Therefore, the focus of the experiments and the data produced came from the Pathmind system, as integrated with AnyLogic. The author will experiment with numerous observations and reward metrics. The reward metrics combine to form the reward function, critical to training in reinforcement learning. The action-space is a rather simple tuple, an array of values containing the state of each dynamic sign. The DRL policy, once trained, gives the AnyLogic model either a TRUE or a FALSE for each dynamic sign. A TRUE means the sign is active and a FALSE that the sign is inactive. This is at the center of the trained policy agent's control over all signs in the scenario. It allows the policy to guide the civilian agents away from the shooter and to a safe exit. Based on the location of each civilian and the location of the shooter, the DRL policy can adjust which signs to activate or de-activate to improve survivability in the modeled scenario. It is critically important to choose observations and set up an effective reward function to allow the DRL policy to learn.

If these metrics are irrelevant to training or are otherwise poorly chosen, the fully trained DRL policy will not yield a desirable result. Adjusting these parameters is the most important aspect of this work. The key to answering the primary research question is observing if the trained policy uses the signs to guide civilians away from the shooter and to a safe exit, thereby reducing casualties. To answer the additional questions the number of signs chosen and where they are placed, was determined by how many times a particular sign was used by civilians to escape. These signs are controlled by the DRL-trained policy; therefore, one can deduce how many signs are needed and where they should be placed. The resulting data was gathered from Monte Carlo runs and analyzed. The chosen observations and reward metrics are listed in Table 3.5 below.

Table 3.5 - Observations and Reward Metrics

Name	**Type**	**Data Type**
locationOfShooter	Observation	double array
locationOfAllCivilians	Observation	double array
signStates	Observation	boolean array
locationOfAllSigns	Observation	double array
locationOfAllExits	Observation	double array
civilianStates	Observation	boolean array
totalCivilianEscaped	Reward Metric	double
totalSignStateChange	Reward Metric	double
currentCasualties	Reward Metric	double

A reward metric is a component of the overall reward function, which is important to reinforcement learning. Pathmind conveniently allows the user to list and import as many of these values as they wish and then combine them into a complete function within the Pathmind web interface. The user may do this simply by indicating in the web interface if a specific metric is to be minimized or maximized during training. Alternatively, should the user choose, he can avoid integrating a metric into the reward function yet still monitor it during training to analyze the training in real-time. The user may code the reward function explicitly, using these metrics or any other values within the interface. All of the observations, reward metrics, and any other values associated with the Pathmind DRL policy are set or coded in the Pathmind Helper within the AnyLogic Model. They may then be edited at will within the Pathmind web interface, as described above. An example of how the observations and reward metrics are coded within the AnyLogic Pathmind Helper is shown below in Figure 3.12.

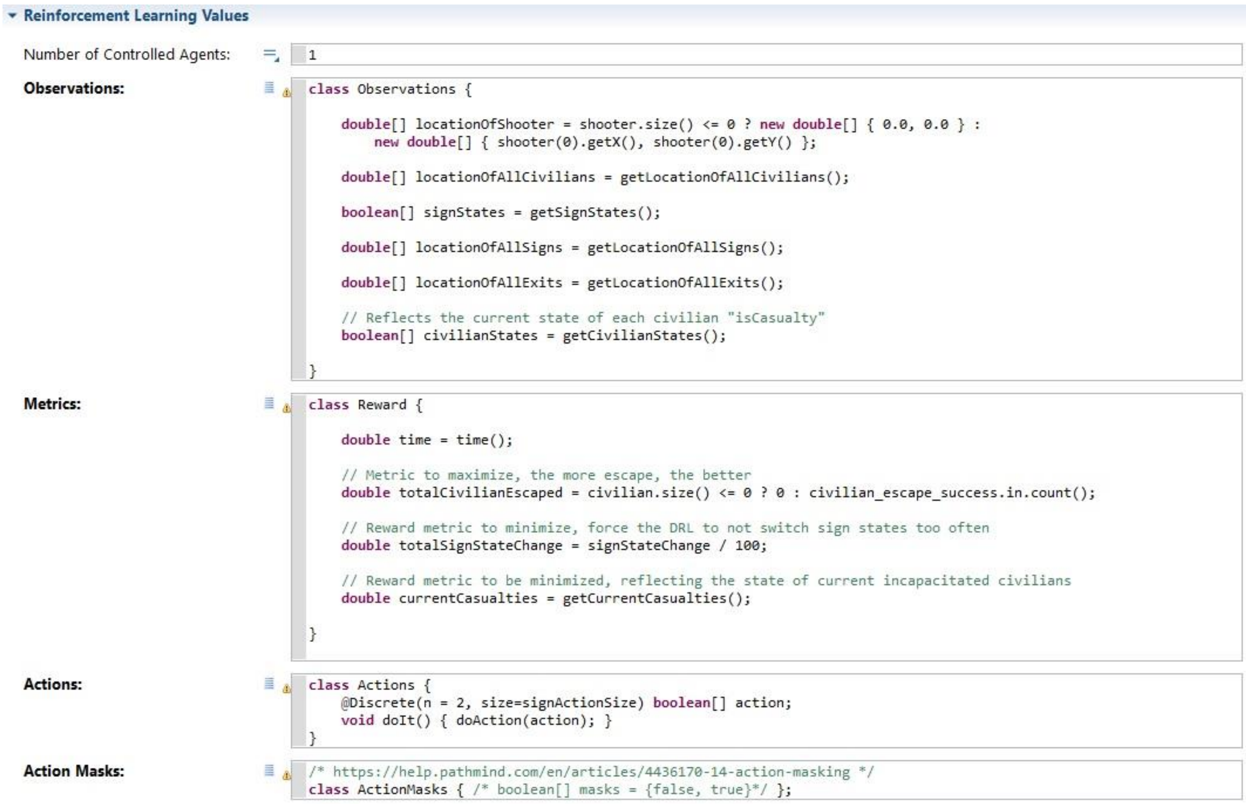

Figure 3.13 - Pathmind Reinforcement Learning Values

The observation and reward metrics listed above were initially chosen by the researcher as most likely to yield positive results. Experimentation was conducted to determine what combination of observations and what reward metrics produced the highest reward after training. These experiments, and the analysis, are explained in detail within Chapter 4.

CHAPTER 4. EXPERIMENTION, ANALYSIS AND RESULTS

4.1 Overview

This chapter outlines the experimental work done to produce data for analysis and establish results. All experimentation is based on the efforts of the previous chapter regarding the model and the deep reinforcement learning (DRL) connection to it. This includes the parameters derived from the foundational model and any new or changed parameters. Three overall sets of experiments were conducted. One, to establish a set of baseline data. Two, experimentation, involving extensive DRL work, helped the author answer the first and overarching research question. Thirdly, experimentation established the foundation to answer both the second and third questions. It is understood that experimentation never truly ends with this type of research and is only halted to report results and meet a specific project scope.

4.2 Baseline Experiments

The purpose of the first set of experiments was to establish baseline data for comparison with the deep reinforcement learning (DRL) policy output. There were two basic scenarios created. The first is to have all civilians pick the closest exit, regardless of any other factor such as the shooter's location, and run towards it. The second included the case of all dynamic signs being activated and allowing the civilians to choose a sign to run to, hopefully making it to an exit. How a civilian picks a sign or an exit was outlined in a previous chapter.

4.2.1 Civilian Run to Closest Exit

A Monte Carlo experiment was used to produce the baseline data for the case of all civilians running to the closest exit. The experiment was set to run 1000 times so as to produce appropriate data to measure. The following data in Table 4.1 was produced for all four scenarios where the difference is in the shooter spawn location only.

Table 4.1 - Run to Closest Exit Results

Shooter Spawn Location	**Mean Runtime**	**Mean Survivors**	**Standard Deviation Survivors**	**Mean Casualties**	**Standard Deviation Casualties**
North	323.18 s	39.31	3.76	60.69	3.76
South	323.43 s	84.99	4.21	15.01	4.21
East	323.26 s	60.73	4.62	39.27	4.62
West	327.84 s	88.89	3.52	11.11	3.52

Of note in the data is that the mean survivors and casualties vary greatly, depending on where the shooter starts at the beginning of the incident. The environment, including obstacles such as structures and walls, affects the shooter reaching the civilians and tends to channelize the civilians when they escape. Of course, the location of where the civilians start at the beginning of the incident, relative to the shooter, plays a role in survival within the model. The author chose to use four separate spawn locations for the shooter for these reasons so as to communicate the results of the research more effectively. Simply using one spawn location might have missed important information.

4.2.2 All Signs Active

The next set of data produced for comparison was a scenario where all signs were set to active. This allowed civilians to detect signs and run towards them, using the algorithm described in a previous chapter. Table 4.2 shows the resulting data output for each shooter spawn location within this set of model parameters.

Table 4.2 - All Signs Active Results

Shooter Spawn Location	Mean Runtime	Mean Survivors	Standard Deviation Survivors	Mean Casualties	Standard Deviation Casualties
North	424.52 s	32.34	3.92	67.66	3.92
South	378.74 s	34.87	4.32	65.13	4.32
East	462.65 s	43.00	5.00	57.00	5.00
West	416.62 s	71.38	5.06	28.62	5.06

The most interesting observation is how the mean casualties are higher than in the previous case of all civilians running to the nearest exit. This is easily understood given that all signs are active, and the civilians have no clear guidance as to which signs lead them away from danger. Whereas in the previous scenario, the civilian agents were directed to run straight to safety, that being an exit. However, in that case, the civilians also had no guidance on the shooter's position; therefore, it was also not an ideal situation for increased civilian survival.

4.3 Pathmind Experimentation

The main effort of this research was the work conducted with the Pathmind deep reinforcement learning (DRL) system. To find the most valuable observations, reward metrics, and neural network combination required extensive experimentation. It was also necessary to regularly modify and tweak the AnyLogic model to accommodate better training of a DRL policy. As the reader will see, the results of the experimentation produced a viable solution to answer all of the research questions. This section gives an overview of the initial work done in this regard, as well as a detailed description of how the author narrowed down the ideal solution for training an effective DRL policy.

4.3.1 Initial Experimentation

Before narrowing in on the end state policy that was used to produce results for analysis and answer the research questions, the author experimented with a wide-ranging set of parameters within Pathmind. Fitting these parameters to the AnyLogic model and eliminating bugs or undesirable behavior of the trained policies was also necessary. The author created 28 separate models within the Pathmind testing environment, each with multiple experiments. The total number of experiments in all of the models uploaded to Pathmind for training was 105. Certainly, not all of these policies trained successfully, yet even those that did not were helpful in discovering bugs and faults within the model and both software systems. All of the experiments led the author to discover what set of observations and reward metrics resulted in the best-trained policy and confirmed that the action space chosen was sufficient. Figure 4.3 demonstrated but a few of the models and their experiments that were run to arrive at the final solution on what combination of parameters would yield the most beneficial answers.

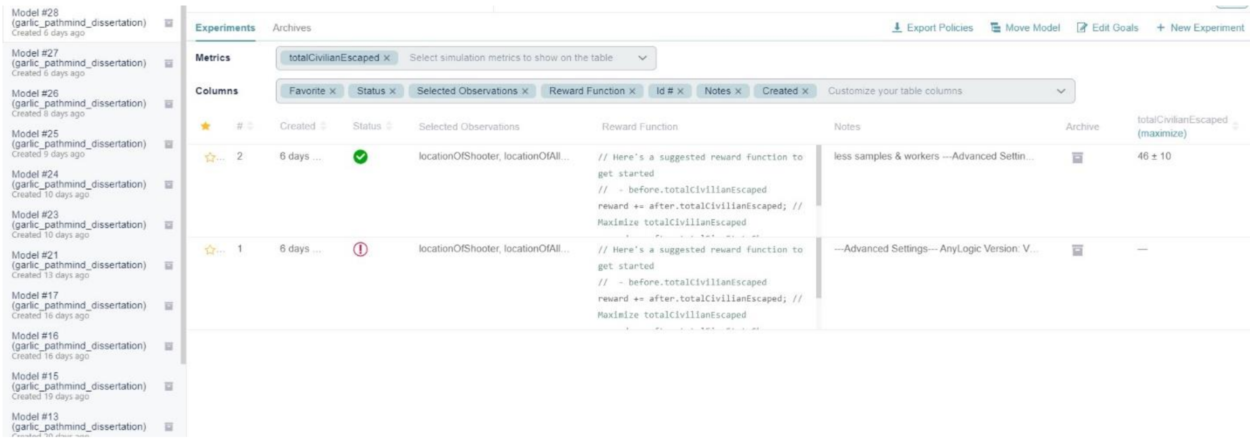

Figure 4.1 - Pathmind Experiments List

Figure 4.1 lists only the active models, including the last one, number 28, that were uploaded for policy training. Other experiments were archived due to various reasons, such as experiment failure or outdated logic in the model. The experiments using all of the 28 models were instrumental in determining the observations and reward metrics used in the final set, which are listed in Table 3.5. In addition to choosing the proper observations and reward metrics, gaining knowledge of the effects of differently configured neural networks was important to train a good policy. The author tried many different neural network configurations with differing depth and width. From shallow networks with only two or three hidden layers, all the way to the maximum allowed in Pathmind, which was ten. He was able to use layers with 64 nodes all the way to 1024 nodes each. The finding was that a deeper network with ten layers, yet with only 64 nodes for each layer, performed best and also reduced training time. Even selecting ten layers with 128 nodes each extended the training time considerably and did not yield a significant improvement in the policy performance. Once the author completed this extensive experimentation, he was able to follow it with the training of the final four policies used to output data for analysis.

4.3.2 Final Experimentation to Select Policy

The author created four experiments for each different shooter spawn location so as to gain one final assessment of the best set of observations and reward metrics to use. This equals a total of 16 experiments. The following set of figures list the four separate sets of observations and reward metrics used in this effort, and the final state of the metrics for each spawn location. It is important to note that the set of four separate parameters used for each set of experiments are the same. That is, each separate spawn location uses the same parameter set for each of the four experiments.

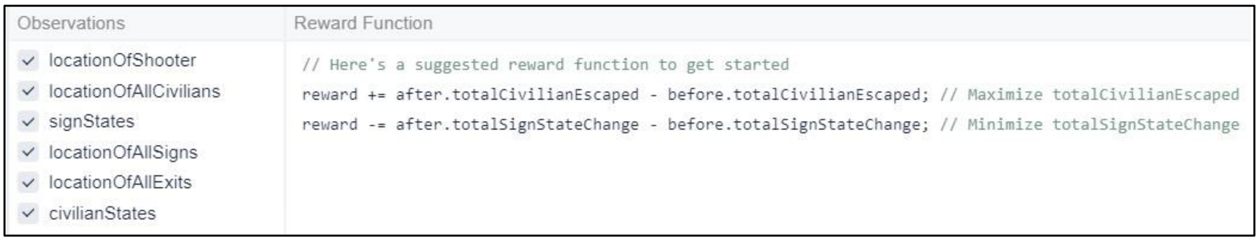

Figure 4.2 - Experiment One Parameters

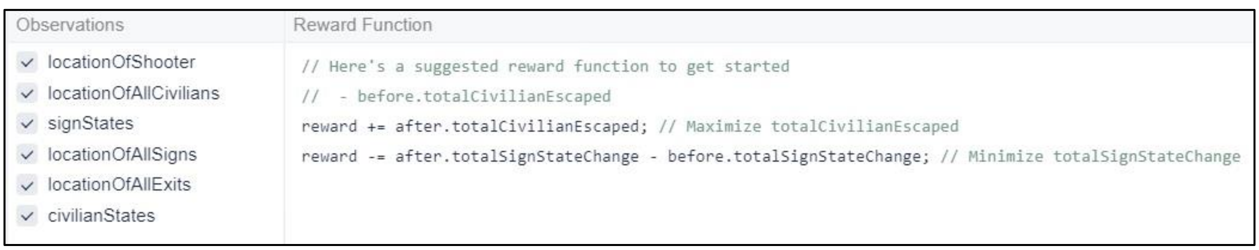

Figure 4.3 - Experiment Two Parameters

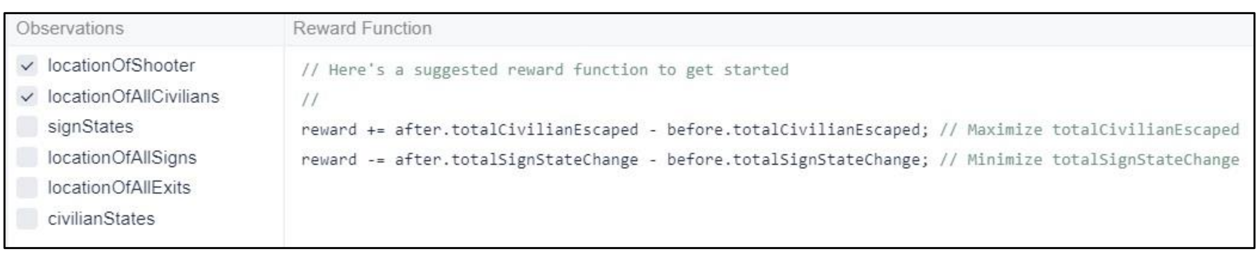

Figure 4.4 - Experiment Three Parameters

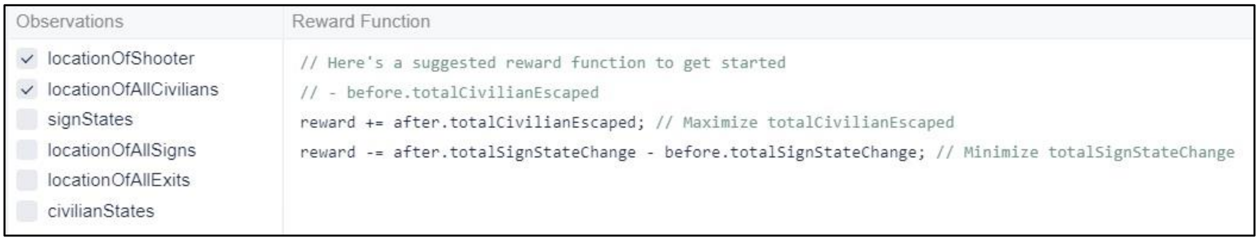

Figure 4.5 - Experiment Four Parameters

Figure 4.6 - North Spawn Experiment One Metrics

Figure 4.7 - North Spawn Experiment Two Metrics

Figure 4.8 - North Spawn Experiment Three Metrics

Figure 4.9 - North Spawn Experiment Four Metrics

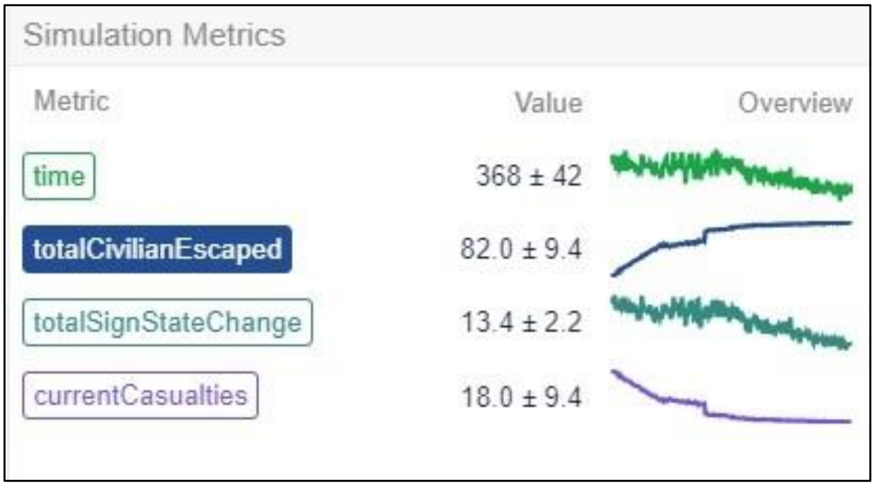

Figure 4.10 - South Spawn Experiment One Metrics

Figure 4.11 - South Spawn Experiment Two Metrics

Figure 4.12 - South Spawn Experiment Three Metrics

Figure 4.13 - South Spawn Experiment Four Metrics

Figure 4.14 - East Spawn Experiment One Metrics

Figure 4.15 - East Spawn Experiment Two Metrics

Figure 4.16 - East Spawn Experiment Three Metrics

Figure 4.17 - East Spawn Experiment Four Metrics

Figure 4.18 - West Spawn Experiment One Metrics

Figure 4.19 - West Spawn Experiment Two Metrics

Figure 4.20 - West Spawn Experiment Three Metrics

Figure 4.21 - West Spawn Experiment Four Metrics

The 16 simulation metrics graphs show each reward metric and other metrics that were useful to monitor during training. The figures represent the final state of training. The observations used for each training experiment are listed in the figures related to parameters. The author chose to select either a set of two or a set of six observations for each training run. This was one of the lessons learned during the previous training events discussed in the initial experimentation section above. Adding or removing individual observations made no significant difference in results, yet having all six or only two moved the needle in either reward metric results or in training time required. This allowed the author to choose between quicker training or better results. The reward function, on the right side of the parameter figures, was composed of only two reward metrics, *totalCivilianEscaped*, and *totalSignStateChange*. The first was designated to be maximized and the other to be minimized. The author intended to maximize the number of survivors. However, he also needed to train the DRL policy to not switch signs from one state to the other too often to avoid having civilians “bounce” between signs. It was discovered during the previous experiments that one reward metric of maximizing escaped civilians did not produce significantly positive results.

The final effort of training 16 policies was useful to confirm the best policy to use for data output and analysis. The most effective policy trained for all spawn locations was that of experiment one, as far as the all-important reward metric of *totalCivilianEscaped* is concerned. This particular set of parameters included all observations and a reward function that subtracts the previous reward from the current reward. This means that at each timestep, the previous reward value is being subtracted from the current, affecting the overall reward for each episode. Each episode is the equivalent to an AnyLogic simulation run and includes multiple steps, depending on how long the episode runs and the Pathmind recurrence. The longer the run takes and the smaller the recurrence number, the more steps there will be. The reward that is generated here after each episode is important for the DRL to learn. This demonstrates that the chosen observation and reward metric combination is the best solution for an effective policy.

4.4 Monte Carlo Experiments with Chosen Policies

Following the selection of the best policies trained in Pathmind for each separate shooter spawn location, Monte Carlo experiments were run within the AnyLogic model, using 1000 runs to output data. These experiments were conducted similarly to the base data discussed previously, yet with different settings for each run. The civilians would now be guided by the trained policy instead of having all signs on to choose from or running straight to the closest exit. The data produced with these experiments is at the center of answering all three research questions.

4.4.1 In Pursuit of Research Question One

The first task after choosing the best deep reinforcement learning (DRL) policies and running Monte Carlo experiments with them was to analyze the data regarding a reduction in civilian agent casualties. This directly supports the goal of answering the first and arguably most important, research question. The following table details the results from the Monte Carlo experiments using the Pathmind trained policies.

Table 4.3 - Pathmind Trained Policy Results

Shooter Spawn Location	Mean Runtime	Mean Survivors	Standard Deviation Survivors	Mean Casualties	Standard Deviation Casualties
North	390.75 s	58.69	5.71	41.31	5.17
South	367.20 s	81.48	4.89	18.52	4.89
East	384.44 s	68.44	4.93	31.56	4.93
West	410.97 s	99.03	2.04	0.97	2.04

To give the reader a better impression of the significance of the results, the author created several figures. Each bar chart represents one different spawn location for each data set and includes the two base cases and the final policy case.

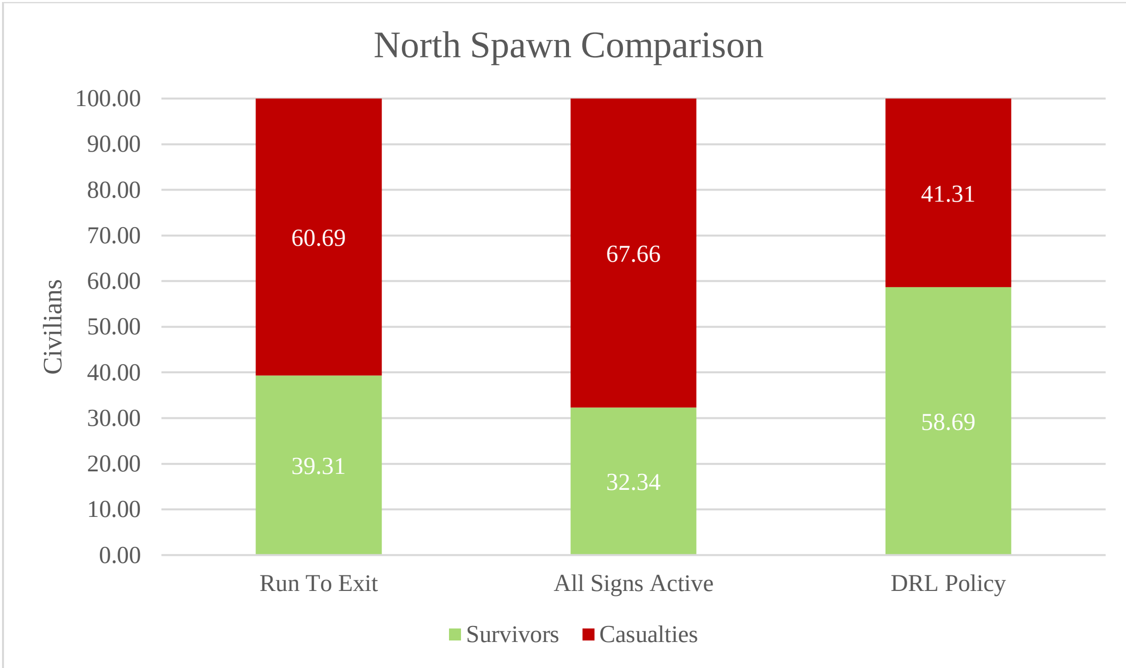

Figure 4.22 - North Spawn Results Comparison

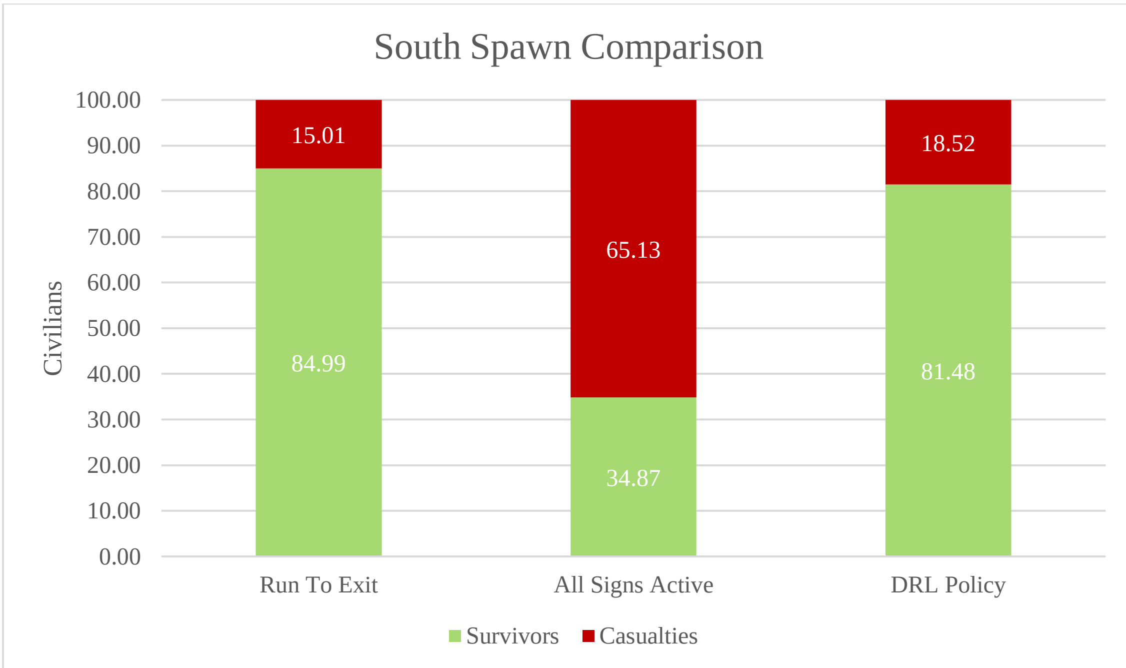

Figure 4.23 - South Spawn Results Comparison

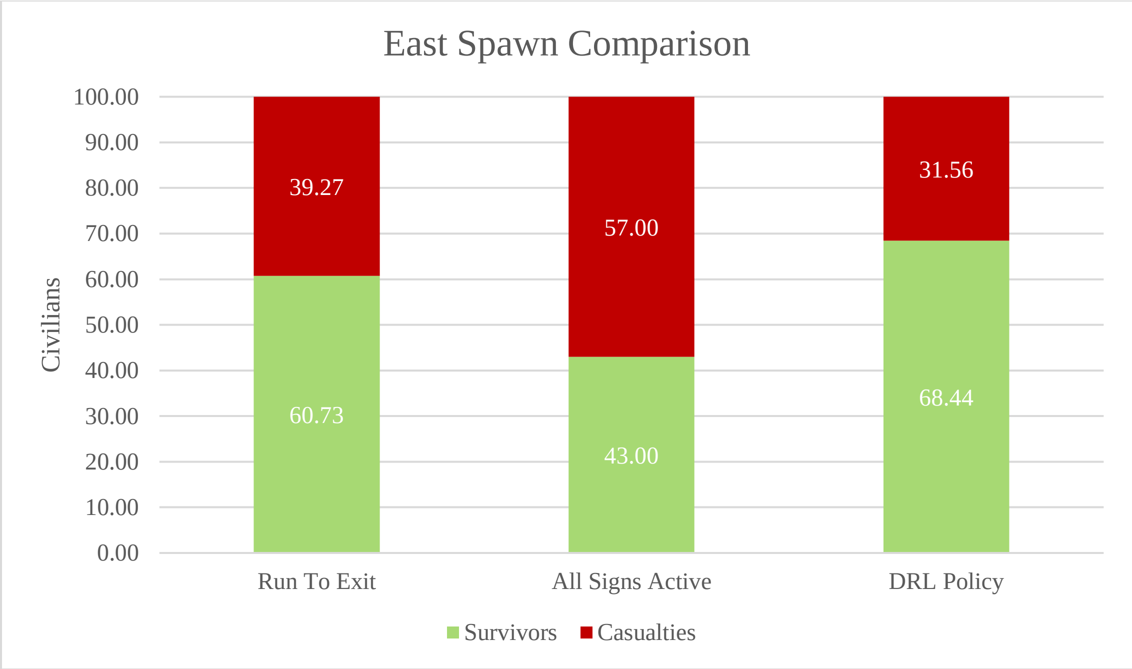

Figure 4.24 - East Spawn Results Comparison

Figure 4.25 - West Spawn Results Comparison

The charts above clearly show that the trained DRL policy performs far and above any of the base cases. More civilian agents survive an incident and reach a safe exit when the DRL policy controls the signs that guide them, hence reducing casualties significantly. The only exception is in the south spawn scenario. There the “run to nearest exit” base case is slightly better than the DRL policy. This shows in part how important the environment is to training proper DRL policies. Given the individual location of civilians at the start of the incident and the position of various obstacles to the shooter in the southern areas of the environment, there is little the DRL policy can do to save more lives. Of course, further experimentation might yield better results, and this will be discussed in the following chapters. Nonetheless, in all other cases, the dynamic signs controlled by the DRL policy reduce casualties significantly compared to the base scenarios. Table 4.4 below is a summary of all results and expresses the difference in casualties as a percentage of casualties relative to the individual spawn cases.

Table 4.4 - Results Comparison Summary

Shooter Spawn Location	**DRL Policy Mean Casualties**	**Run To Exit Mean Casualties**	**DRL Policy Difference Percentage**	**All Signs Active Mean Casualties**	**DRL Policy Difference Percentage**
North	41.31	60.69	-31.92 %	67.66	-38.94 %
South	18.52	15.01	23.35 %	65.13	-71.56 %
East	31.56	39.27	-19.63 %	57.00	-44.62 %
West	0.97	11.11	-91.30 %	28.62	-96.62 %

The table demonstrates the effectiveness of a DRL-trained policy in reducing casualties. In all but one comparison, there is a significant reduction in casualties, leading the author to believe that this technology has promise in the field of active shooting incidents (ASI’s).

4.4.2 In Pursuit of Questions Two and Three

Further analysis of the data collected via the trained DRL polices answered the remaining research questions. In addition to the data discussed above, information on how many signs were used by civilians across 1000 Monte Carlo runs for each shooter spawn location case was collected. Every time a civilian agent detected a sign and chose to run towards it within each Monte Carlo run, the number for this sign was incremented. This provided a good understanding of what signs at what locations were important to evacuate civilians. If a sign was never used by even one civilian in 1000 runs, one can safely assume it is unimportant within the Garlic Festival environment. These signs would not contribute to increasing the survivability of civilians and decreasing casualties. This data forms the basis for the analysis needed to answer research questions two and three.

The data below within Table 4.5 lists all signs placed in the Garlic Festival environment within the model before any reduction occurred. There are 119 signs total spread across the terrain. Each sign is 60 feet apart within a simple matrix, and their coordinates X and Y represent the location within the AnyLogic canvas, starting in the upper left corner. The X coordinate represents the East-West direction, and the Y coordinates the North-South direction are measured in pixels within AnyLogic. The pixel distances translate to feet within the scale of the environment. It is important to note that some signs (called attractors in the system) were placed manually for reasons discussed in Chapter 3.

Table 4.5 - All Signs and Locations Pre-Reduction

Sign Name	**X**	**Y**
attractor	24.54	24.54
attractor1	73.621	24.54
attractor2	122.702	24.54
attractor3	171.782	24.54
attractor4	220.863	24.54
attractor5	269.944	24.54
attractor6	319.025	24.54
attractor7	368.105	24.54
attractor8	417.186	24.54
attractor9	466.267	24.54
attractor10	515.347	24.54
attractor11	564.428	24.54
attractor12	613.509	24.54
attractor13	24.54	73.621
attractor14	73.621	73.621
attractor15	122.702	73.621
attractor16	171.782	73.621
attractor17	220.863	73.621
attractor19	316.025	73.621
attractor20	368.105	73.621
attractor21	417.186	73.621
attractor22	466.267	73.621
attractor23	515.347	73.621
attractor24	564.428	73.621
attractor25	613.509	73.621
attractor26	662.589	73.621
attractor27	711.67	73.621
attractor28	24.54	122.702
attractor29	73.621	122.702
attractor30	122.702	122.702

Table 4.5 continued

attractor31	171.782	122.702
attractor32	220.863	122.702
attractor34	321.025	122.702
attractor35	368	123
attractor36	417.186	122.702
attractor38	515.347	122.702
attractor39	564.428	122.702
attractor40	613.509	122.702
attractor41	662.589	122.702
attractor42	711.67	122.702
attractor43	760.751	122.702
attractor44	24.54	171.782
attractor46	122.702	171.782
attractor47	171.782	171.782
attractor48	220.863	171.782
attractor49	269.944	171.782
attractor50	319.025	171.782
attractor52	417.186	171.782
attractor53	466.267	171.782
attractor54	515.347	171.782
attractor55	613	169
attractor57	662.589	171.782
attractor58	711.67	171.782
attractor59	760.751	171.782
attractor60	24.54	220.863
attractor61	68.621	220.863
attractor63	171.782	220.863
attractor66	319.025	220.863
attractor67	368.105	220.863
attractor68	417.186	220.863
attractor69	466.267	218.863

Table 4.5 continued

attractor70	515.347	220.863
attractor71	564.428	220.863
attractor72	613.509	220.863
attractor73	662.589	220.863
attractor74	711.67	220.863
attractor75	24.54	269.944
attractor76	73.621	269.944
attractor77	122.702	269.944
attractor78	171.782	269.944
attractor79	220.863	269.944
attractor81	318.025	269.944
attractor82	368.105	269.944
attractor84	466.267	269.944
attractor86	565	270
attractor88	662.589	269.944
attractor89	24.54	319.025
attractor90	73.621	319.025
attractor91	122.702	319.025
attractor92	171.782	319.025
attractor93	220.863	319.025
attractor94	269.944	319.025
attractor95	319.025	319.025
attractor96	368.105	319.025
attractor97	417.186	319.025
attractor98	466.267	319.025
attractor99	515.347	319.025
attractor101	613.509	319.025
attractor102	73.621	368.105
attractor103	122.702	368.105
attractor104	171.782	368.105
attractor105	220.863	368.105

Table 4.5 continued

attractor106	269.944	368.105
attractor107	319.025	368.105
attractor108	368.105	368.105
attractor109	420	368
attractor110	466.267	368.105
attractor111	515.347	368.105
attractor112	564.428	368.105
attractor113	613.509	368.105
attractor114	220.863	417.186
attractor115	269.944	417.186
attractor116	319.025	417.186
attractor117	368.105	417.186
attractor118	417.186	417.186
attractor119	466.267	417.186
attractor120	269.944	466.267
attractor122	515	418
attractor123	417.186	466.267
attractor124	466.267	466.267
attractor125	319.025	515.347
attractor126	368.105	515.347
attractor127	417.186	515.347
attractor128	369	550
attractor129	417	550
attractor130	800	145
attractor131	760	220
attractor132	465	516
attractor133	465	550

The results of an analysis of the sign data produced by Monte Carlo run, for each spawn location case, is shown and discussed below. The results produced are at the core of answering the two final research questions.

For the north spawn location case, only 95 signs were used by the DRL policy to achieve the results discussed in Section 4.4.1 regarding civilian survival. This means that 24 signs can be removed from the environment giving their irrelevancy in affecting any change in civilian survival. The signs that were removed are listed below in Table 4.6.

Table 4.6 - North Spawn Case Removed Signs

Sign Name	**X**	**Y**
attractor	24.54	24.54
attractor1	73.621	24.54
attractor2	122.702	24.54
attractor3	171.782	24.54
attractor4	220.863	24.54
attractor13	24.54	73.621
attractor14	73.621	73.621
attractor15	122.702	73.621
attractor16	171.782	73.621
attractor28	24.54	122.702
attractor29	73.621	122.702
attractor30	122.702	122.702
attractor32	220.863	122.702
attractor35	368	123
attractor44	24.54	171.782
attractor46	122.702	171.782
attractor60	24.54	220.863
attractor61	68.621	220.863
attractor75	24.54	269.944
attractor76	73.621	269.944
attractor89	24.54	319.025
attractor90	73.621	319.025
attractor96	368.105	319.025
attractor102	73.621	368.105

For the south spawn case, only 92 signs remain after removing the unnecessary ones from the environment, meaning 27 can be removed. Table 4.7 below lists all of the removed signs.

Table 4.7 - South Spawn Case Removed Signs

Sign Name	**X**	**Y**
attractor	24.54	24.54
attractor1	73.621	24.54
attractor2	122.702	24.54
attractor3	171.782	24.54
attractor4	220.863	24.54
attractor5	269.944	24.54
attractor13	24.54	73.621
attractor14	73.621	73.621
attractor15	122.702	73.621
attractor16	171.782	73.621
attractor17	220.863	73.621
attractor20	368.105	73.621
attractor28	24.54	122.702
attractor29	73.621	122.702
attractor30	122.702	122.702
attractor31	171.782	122.702
attractor35	368	123
attractor44	24.54	171.782
attractor46	122.702	171.782
attractor60	24.54	220.863
attractor61	68.621	220.863
attractor75	24.54	269.944
attractor76	73.621	269.944
attractor89	24.54	319.025
attractor90	73.621	319.025
attractor102	73.621	368.105
attractor125	319.025	515.347

For the east spawn case, only 96 signs remain after removing the unnecessary ones from the environment, meaning 23 can be removed. Table 4.8 below lists all of the removed signs.

Table 4.8 - East Spawn Case Removed Signs

Sign Name	**X**	**Y**
attractor	24.54	24.54
attractor1	73.621	24.54
attractor2	122.702	24.54
attractor3	171.782	24.54
attractor4	220.863	24.54
attractor5	269.944	24.54
attractor6	319.025	24.54
attractor13	24.54	73.621
attractor14	73.621	73.621
attractor15	122.702	73.621
attractor16	171.782	73.621
attractor28	24.54	122.702
attractor29	73.621	122.702
attractor30	122.702	122.702
attractor44	24.54	171.782
attractor46	122.702	171.782
attractor60	24.54	220.863
attractor61	68.621	220.863
attractor75	24.54	269.944
attractor76	73.621	269.944
attractor89	24.54	319.025
attractor90	73.621	319.025
attractor102	73.621	368.105

For the west spawn case, only 90 signs remain after removing the unnecessary ones from the environment, meaning 29 can be removed. Table 4.9 below lists all of the removed signs.

Table 4.9 - West Spawn Case Removed Signs

Sign Name	X	Y
attractor	24.54	24.54
attractor1	73.621	24.54
attractor2	122.702	24.54
attractor3	171.782	24.54
attractor4	220.863	24.54
attractor5	269.944	24.54
attractor6	319.025	24.54
attractor7	368.105	24.54
attractor13	24.54	73.621
attractor14	73.621	73.621
attractor15	122.702	73.621
attractor16	171.782	73.621
attractor17	220.863	73.621
attractor19	316.025	73.621
attractor20	368.105	73.621
attractor28	24.54	122.702
attractor29	73.621	122.702
attractor30	122.702	122.702
attractor31	171.782	122.702
attractor44	24.54	171.782
attractor46	122.702	171.782
attractor60	24.54	220.863
attractor61	68.621	220.863
attractor75	24.54	269.944
attractor76	73.621	269.944
attractor89	24.54	319.025
attractor90	73.621	319.025
attractor102	73.621	368.105
attractor125	319.025	515.347

It is easier to understand and imagine what signs were removed and to analyze the results with proper figures. Figure 4.26 below is the default case with all signs in the environment, placed here again for the convenience of the reader to compare. Figures 4.27 through 4.30 are the north, south, east, and west spawn location cases, where the signs in the tables above were removed. All signs are represented as AnyLogic attractors, and the signs being used in each case are highlighted in magenta color. The black attractors are the ones removed from the respective scenarios.

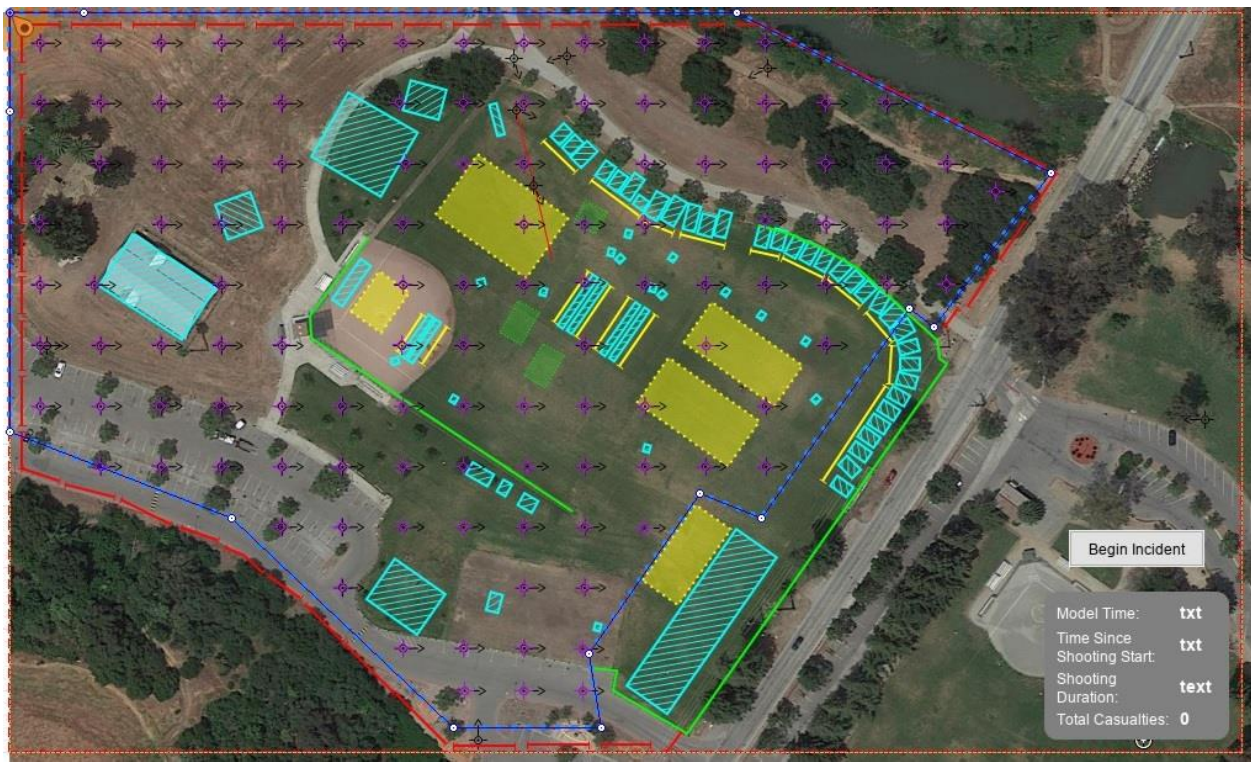

Figure 4.26 - Base Sign Placement Before Individual Removal

Figure 4.27 - North Spawn Case Removed Signs

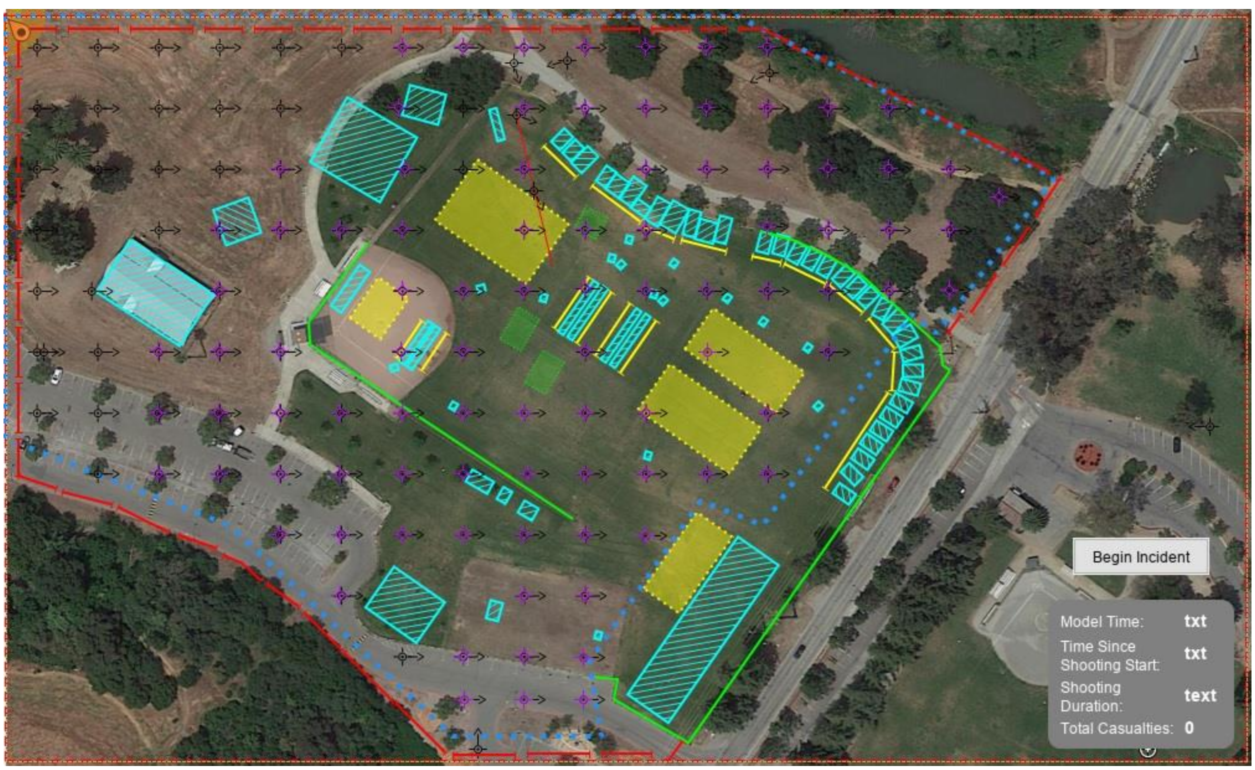

Figure 4.28 - South Spawn Case Removed Signs

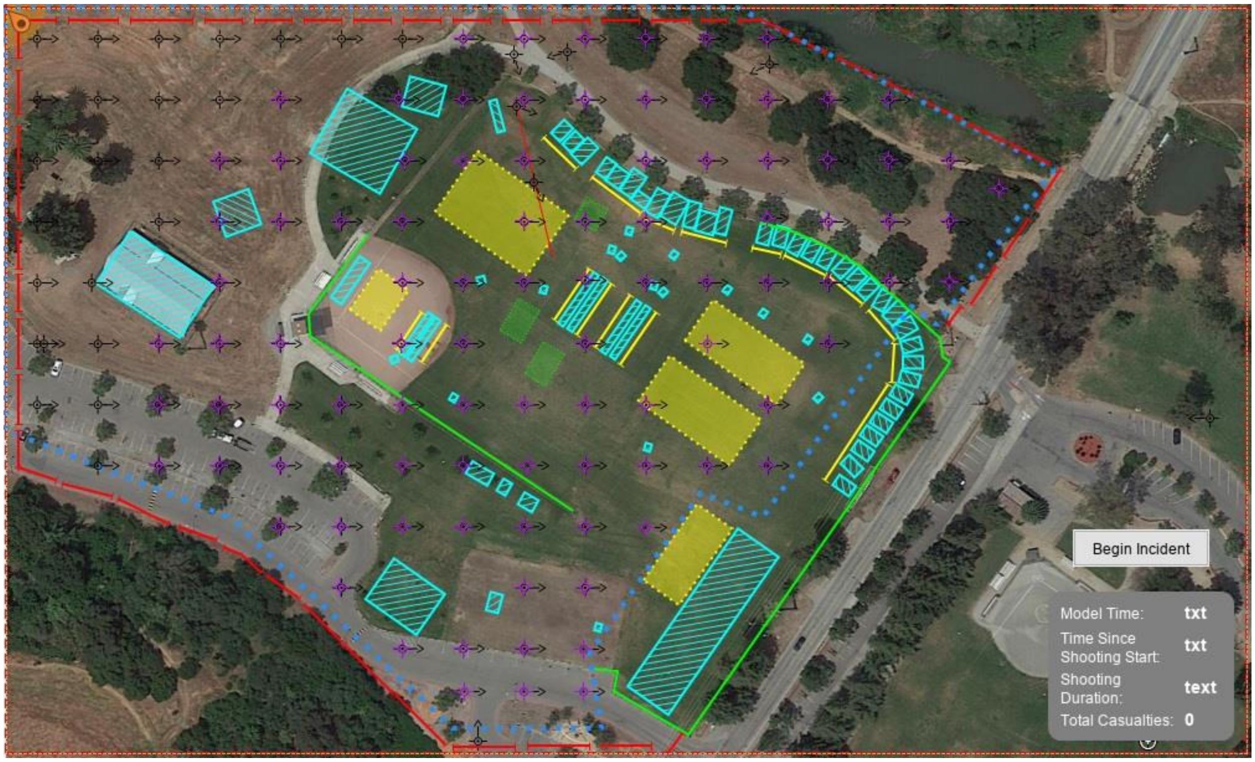

Figure 4.29 - East Spawn Case Removed Signs

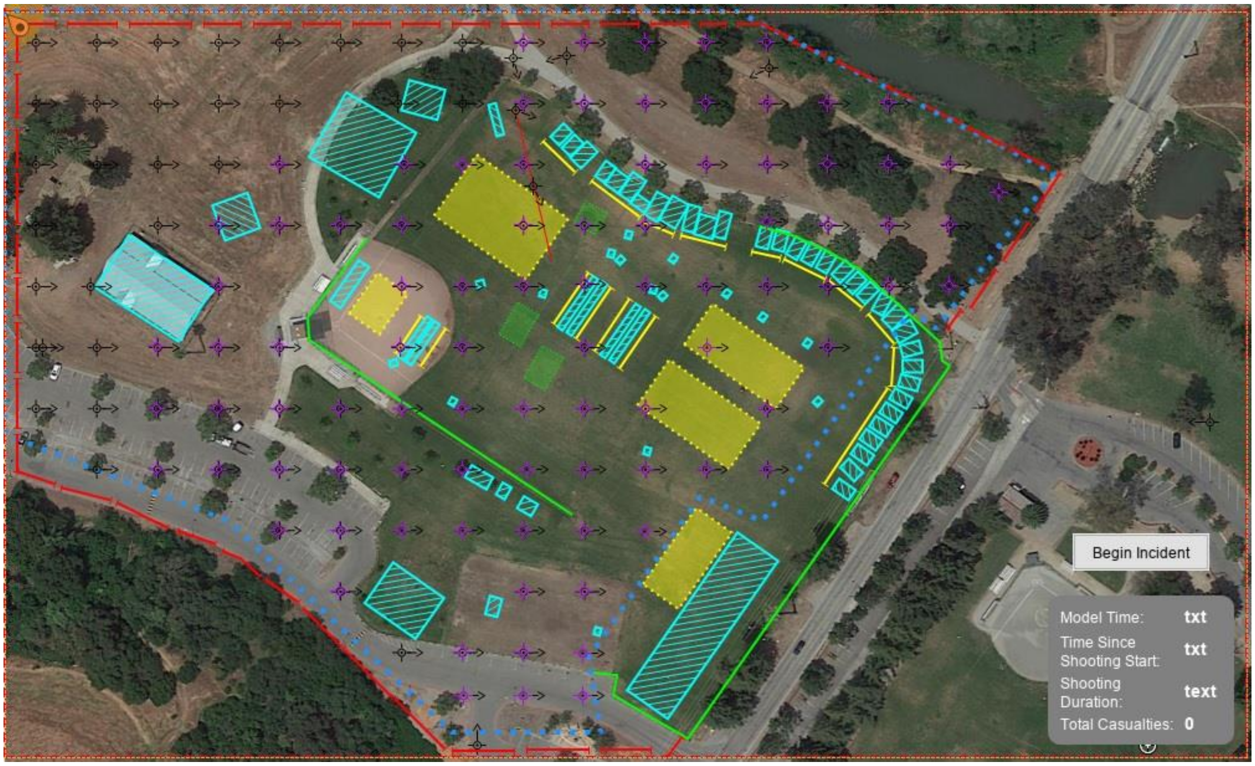

Figure 4.30 - West Spawn Case Removed Signs

Comparing the images, one can see clearly that the signs placed on the west side of the environment are unimportant and that there is a clear preference of the DRL policy to use signs close to and around the yellow civilian gathering areas. Also, no matter where the shooter spawns, there is a clear priority given to evacuating towards the north or the south. This is most likely due to the proximity of the yellow-colored civilian gathering areas to those two general directions; that is, the majority of the civilians will congregate north and south. Further analysis and re-training of the DRL policy with these reduced-sign environments could yield a more precise reduction of signs needed. The author did re-train the policies without the signs, yet stopped at one iteration of the process, since the purpose of this research is not to find the most optimal number and location of signs, but only if DRL can help us find those answers. Further refinement work and what it might yield will be discussed in a later chapter on future research.

4.4.3 Confirming Reduced Sign Scenarios

To be confident in the state of the DRL trained policy with the changed environments, the author re-trained each policy for each separate spawn case with the reduced signage. It is important to confirm that the results are still advantageous to civilian survival so that the answers to the final research questions are validated. Though, even without validation, we can assume that the DRL policy is telling us that certain signs are not needed. However, given that the environment has changed and that DRL has a critical dependence on it, not retraining policies in this new situation could be a fallacy. Therefore, the author took the time to train these new policies and produce the data that is presented below in table 4.10.

Table 4.10 - Reduced Sign Casualty Comparison

Shooter Spawn Location	Maximum Number of Signs Casualties	Reduced Signs Casualties
North	41.31	44.44
South	18.52	16.13
East	31.56	31.78
West	0.97	5.90

The number of casualties in the reduced sign scenarios is similar to the scenarios using the maximum number of 119 signs. Although, the west spawn scenario has a rather large increase in casualties, relative to the original number. Given that most of the signs were removed from the west of the environment, further study into these results is warranted. This further exploration is something to consider for future research using this model.

4.5 Research Question One

What impact does the application of reinforcement learning have on the number of casualties during an active shooting incident?

4.5.1 Answering Research Question One

The impact on applying reinforcement learning (RL) to the problem of active shooting incidents (ASIs) is clear. There is a reduction of casualties in the Garlic Festival environment when deep reinforcement learning (DRL) trained policies are in control of dynamic signs that control civilian agent movement. In each and every scenario, save one, there is a significant reduction of lives lost when applying these technologies together. In particular, when compared to the "all signs active" case, the DRL policy yields the most consistent results. This comparison is most valid.

This is because, in a real-world incident, confusion, fear, and terror reigns in the minds of civilians, and this is best simulated by having all signs on without clear guidance to the civilian agents. Therefore, comparing this to the DRL-trained policy is most appropriate. We see an average reduction of casualties of 62.94% across all spawn location cases. Even with the more advantageous and less realistic case of all civilians knowing where the closest exit is and running to it immediately, we see an average reduction of casualties of 29.88% across all four spawn location cases, using the DRL policy.

These results suggest a substantial positive impact in using DRL to seek solutions in the ASI research field. As with other fields in private industry, government and research discussed above, researchers and professionals seeking answers to improving survivability during ASI's would be wise to consider DRL technology.

4.6 Research Question Two

How many reinforcement learning controlled dynamic signs are needed to reduce casualties within an environment?

4.6.1 Answering Research Question Two

The number of signs needed varies depending on the specific environment modeled. Included in "the environment" are things like structures, obstacles, and the position of dynamically changing objects, such as civilian agents and the shooter. It is inherent in deep reinforcement learning (DRL) to keep the environment constant to properly change. By constant, the author means not introducing new variables after training the policy since the DRL policy would not be able to generalize on unknowns it has not seen during training. For example, the policy that was trained with a north spawning shooter would not be as effective if the shooter spawned in the south.

During training, the DRL algorithm would not have necessarily been exposed to many states generated by a shooter appearing in a different location as well as the various states of civilians reacting to the shooter.

It was found through experimentation that the number of signs needed in the four separate environments was as follows. For the north spawn case, a minimum of 95 signs was required; for south spawn, 92 signs, 96 signs for east, and 90 for the west spawning shooter case. The author wishes to point out that further refinement could be conducted to reduce the number of needed signs even further. He will discuss this in the following chapters. For the purpose of this study, it is objectively clear that the DRL policy can and has assisted in determining how many signs are needed to reduce casualties in the four respective environments.

4.7 Research Question Three

At what locations should reinforcement learning controlled dynamic signs be placed to reduce casualties within an environment?

4.7.1 Answering Research Question Three

The third question is directly tied to question number two. The exact locations of signs needed in the four separate environments are listed within tables 4.5 through 4.9. The exact coordinates of each remaining sign for each shooter spawn case is listed therein. Also, figures 4.27 through 4.29 provide a good understanding of where each sign must be located within the Garlic Festival environment. The deep reinforcement learning (DRL) policy has shown us clearly where the signs must be to reduce casualties. Just as with question two, further refinement can be done, which will be discussed in the following chapters.

4.8 Further Experimentation

The author decided to round out his work with an experiment to test if a deep reinforcement learning (DRL) policy trained in a different environment would still perform satisfactorily or fail to reduce casualties. He chose to run a Monte Carlo experiment with 1000 runs using the northern spawn case trained policy with a southern spawn set parameter. As indicated above, the policy trained in one environment should not perform well when used in another environment. Table 4.11 below summarizes the results from this experiment.

Table 4.11 - Policy Comparison in Differing Environment

Shooter Spawn Location	Mean Runtime	Mean Survivors	Standard Deviation Survivors	Mean Casualties	Standard Deviation Casualties
North	390.75 s	58.69	5.71	41.31	5.17
South	367.20 s	81.48	4.89	18.52	4.89
South with North Spawn Policy	380.73 s	46.14	6.89	53.86	6.89

The policy performs poorly compared to the one trained in the same environment. The northern spawn trained policy applied to the southern spawn case only produces a mean of 46.14 survivors, vs. 81.48 survivors with the southern trained policy intended for this environment. However, compared to one of the base scenarios, with all signs turned on, even this misguided policy performs better. In the southern spawn case, where all signs are on, providing no guidance to fleeing civilian agents, the mean survivors were 34.87. Therefore, the misappropriated policy still saves more lives than one of the base scenarios. This is a positive finding since it might indicate that even in an unknown environment a DRL trained policy can still help increase overall survivability. This is a good prospect for future research and application of this technology in real-world environments.

CHAPTER 5. SUMMARY AND RECOMMENDATIONS

5.1 Overview

There has been a steady and observable increase in active shooting incidents (ASIs) over the last 20 years within the United States of America. The majority of these shootings occur in closed spaces such as schools, government buildings, and businesses. There is, however, a significant number that occurs in open-air venues and open spaces, and at least one of these has been particularly lethal (Corcoran et al., 2019). Researchers should continue to support decision makers with useful and actionable studies allowing for the development of effective policy to guard against such violence.

This study was imagined and produced to support the creation of policies that will help defend innocents at the point of friction, or “in extremis”. It is not intended to contribute to policies that help prevent the occurrence of ASIs. The author thought to use the latest developments in technology and tools to gain better insight into ways to handle the defense against active shooters. The environment was explicitly chosen to experiment in an open-air venue so as to test the efficacy of dynamic signs combined with deep reinforcement learning (DRL) control. As Frantz (2021, p. 107) rightfully states, “The recent increase in open-air active shooter events demands investigation and testing of non-traditional methods to reduce casualties.”. The author considers his work to be within this “non-traditional” method, given his investigation into dynamic signs and their use with DRL. Either way, the envelope must be pushed in this field to ensure the increased survival of innocents during these violent incidents.

5.2 Significance of this Study

Given the steady increase in active shooting incidents (ASIs), more research should be devoted to this subject in support of policy makers. Much has been done in the years after the Columbine shooting to further all aspects of defense, with the intention of reducing casualties and bringing the incident to a speedy end. The majority of this research has been focused on law enforcement responder tactics and training, and to a lesser degree, the actions of the civilians caught in the mayhem. This research and the resulting policies have already saved lives and continue to be validated during and after each tragic event.

However, not nearly as much effort and funding has been invested into integrating the latest technologies to help responders and civilians during an incident. There has been a distinct lack of focus on ASIs perpetrated in open-air venues and open spaces. The author’s work is a significant step toward integrating technology and its use in open-air venues and open spaces by testing yet unproven concepts and ideas within a simulated environment. This work demonstrates clearly that agent-based modeling and simulation (ABMS) is a valuable tool to this end. New technologies such as deep reinforcement learning (DRL) can also be successfully integrated into a model to yield useful and actionable results. This study showed that the average reduction in casualties was 62.94% and 29.88%, in comparison to two separate scenarios that did not use DRL-controlled dynamic signs.

With these results, policy makers can make decisions in support of new technologies and fund further research and development. From the agent-based modeling to the prototyping of dynamic signs and their control via machine learning (ML), any effort spent on these new approaches will yield positive real-world results. These technologies are not new and are used in many other fields of research and in business, giving further confidence to this approach. It is the application of these technologies to the active shooting research space that is most significant.

Using technology like this during an incident can and will assist all involved to survive, and as Frantz (2021, p. 108) states, “Dynamic signage that can be adjusted real-time would benefit outdoor venues with limited areas for people to take cover during an active shooter event.”. Other technologies can also be integrated with dynamic signage, which is discussed below.

5.3 Future Technology Development

It would behoove policy makers to fund research and development into various technologies to defend against active shooters. One such technology, dynamic signage, integrated with deep reinforcement learning trained policies, has been discussed in this work. Another well-developed technology is the unmanned aerial vehicle (UAV), commonly referred to as “drones”. These flying objects could help detect, deter and defeat active shooters in short time if properly employed. Acoustic and visual sensors could also be used to help detect and pinpoint a shooter’s location, helping responders, including UAVs, intercept the shooters quickly. The sooner the incident ends, the fewer innocents will be harmed. As this research and the resulting model abstract the difficult task of detecting a shooter and his location, the author will propose several solutions to accomplish this in the future, including research suggestions related to the current model to help develop these technologies more quickly.

5.3.1 Dynamic Signage

The dissertation upon which parts of this work are based, listed a patent for dynamic signage that was a catalyst for the current research. Frantz (2021, p. 109) states that “Implementation of dynamic signage, or signage that can be adjusted real-time, is a defense mechanism an open-air venue can employ to protect patrons in the event of an emergency situation.”. Given the confusion and panic during a life-or-death event such as an active shooting incident (ASI), a highly visible sign that can help people evacuate in a safe direction is worth the investment. There is potential to save many lives. Further research and development should be conducted to build prototypes and integrate the deep reinforcement learning (DRL) trained policy into the system of control. The details of the current state of development and the patent involving such signage is explained by the previous researcher (Frantz, 2021). Further development is required on the DRL policy within the simulated environment, as well as the changing of the fidelity of the model used in training.

5.3.2 Unmanned Aerial Vehicles

The availability and sales of low-cost and commercially off-the-shelf (COTS) UAVs have grown exponentially over the last few years (Saracco, 2019). As of 2019, there are approximately 1.1 million commercial UAVs within the United States, and that number is expected to rise to 3.5 million by 2021 (Fleming, 2019). These "drones" are already used by some police departments and sheriff's offices to help in law enforcement, including during an active shooting. During the ASI in 2021 at a supermarket in Boulder, Colorado, these UAVs were used to monitor the situation and detect the location of the shooter, who later barricaded himself in the building (Rubino, 2021). The increase in production and use of UAVs has considerably reduced their cost to purchase and operate. UAV cost-benefit improvement is a trend that will continue in the coming years. This makes the possibility of using mini and micro-UAVs for a fully integrated system to detect, deter and defeat a shooter while directing responders to the scene and evacuating civilians viable. Together with the integration of DRL and other machine learning (ML) technologies, a small UAV can be of value in defeating an active shooter in extremis. The author proposes to not only use these systems for observation, but also for more offensive actions. Consider a system that is semi or fully autonomous, with cheap yet effective sensors to detect a shooter, that dispatches these

UAVs immediately upon detection of the threat. This autonomous system would seek out and confront the shooter within an open-air venue. A UAV equipped with a non-lethal weapon can at the very least suppress the shooter quickly so as to give civilians time to run and responders time to confront and neutralize the perpetrator. A UAV equipped with irritants such as pepper spray or pepper balls can significantly impede the shooter so as to render him temporarily blind and ineffective. It is clear from literature and research that time is the most important variable in reducing casualties during an ASI; each second might be a life lost or changed forever. Of course, systems like this need to have a "man in the loop" to ensure that certain boundaries are not crossed by the machine-controlled UAV. For the sake of future research and development, however, the prospect of intercepting a shooter quickly with such technology offers immense potential in ASI mitigation.

5.3.3 Sensors for Detecting and Locating a Shooter

Detecting a shooter quickly after he begins his rampage is a difficult task. During the chaos and confusion, particularly within the early minutes, it is hard for responders to fix on the shooter's position, such as the Parkland shooting incident. If certain technologies that help in detection could be leveraged and deployed in open-air venues and open spaces, it would significantly reduce the time a shooter has to inflict damage. Systems that use acoustics or visual sensors might, with the help of ML, assist in directing assets to the shooter's location, including autonomous UAVs. One researcher discussed the use of gunshot triangulation systems, such as the ShotSpotter system (Frantz, 2021). Similar to UAVs, the cost of these systems have decreased significantly over the years, as Frantz (2021) points out in his work. This allows more law enforcement agencies to purchase these systems and to install them around larger areas. According to Frantz (2021, p. 38), the system's "sensor network has the ability to decipher between single or multi-shot weapon systems based upon the acoustic profiles, which is important information to share with first responders”. An integrated system can then make an initial determination of where the shooter is and what weapons he might be using, starting the process of interception by responders and/or autonomous systems. Visual sensors could also contribute to the task of detecting and pinpointing a shooter. Similar to how UAVs and acoustic systems have become more affordable, high-resolution cameras have also become less cost prohibitive, and their widespread availability has only increased. Military application of visual sensors has continued to develop, and many of these technologies can be adopted by law enforcement. Infrared (IR) technology is one such solution that adds the capability to detect a shooter at the onset of an incident. IR capabilities are available for civilian and law enforcement acquisition and are also becoming more affordable. Together with ML, specifically supervised learning, a camera could easily detect a gunshot, even on a bright and hot day. At the same time, an array of these affordable cameras can tag and continuously track a shooter from the moment of identification until the inevitable interception by responders. Together with acoustic triangulation, the system would have a high probability of confirming a shooting has begun and alerting other autonomous systems and human responders.

5.3.4 Integrated Defensive System

The author proposes to combine the technologies discussed above into a fully integrated system of defense similar to the military’s Integrated Air Defense System (IADS), which exists to counter a multitude of threats from the air (Mattes, 2019). It combines multiple systems such as low and high-altitude missiles, anti-air artillery, and a variety of sensors. It is guided by well-established procedures to operate all assets and includes computerized command and control systems. These systems combine to form a strong defense against a specific aerial threat and can act as a model to do the same against active shooters and other threats in open-air venues.

The fundamental idea is to combine optical and acoustic sensors for detection with autonomous UAVs for immediate response and dynamic signage for evacuation. This will then, of course, be added to the traditional security measures already in place. If a shooting can be detected and pinpointed quickly, it will save lives. An immediate response from a UAV armed with non-lethal weapons can suppress a shooter, buying even more time for law enforcement to respond. At a minimum, the shooter can be tracked by the UAV feeding critical information to responding security. This interception would take place while civilians are guided away from the threat using dynamic signs. Practically speaking, much research and development is required to make such a product viable from a monetary perspective. However, given some of the injuries and deaths from recent shootings in open spaces, it might be justified from a risk management perspective, adding impetus for funding. The system could be fielded by private contracting companies that provide it as a per-event service, together with the more traditional security consultation. Alternatively, it might be purchased and run by local, state, or federal law enforcement agencies to provide security at large events regionally and nationally.

This system could also be trained and configured to detect other forms of crime, even theft and mob violence that often occur at large venues. Similar to the IADS, this integrated defensive system could observe an area and watch for a wide array of incidents and alert authorities or use its autonomous systems to counter the threats.

5.4 Future Research

This research was focused on using deep reinforcement learning (DRL) to help evacuate civilians away from a threat. It was determined that the application of this technology in the active shooting incident (ASI) research field is legitimate and can yield positive results. However, more work must occur to improve the specific model used in this research, including high and low-fidelity components, to find higher applicability to the real-world environment.

5.4.1 Improving the Current DRL Policy

First it is useful to discuss what might improve the trained policy and yield a higher reward score, which translates to a higher number of survivors. It helps to illustrate what other observations fed to the DRL neural network might help push the score higher, as well as what reward metrics might improve the end result. Suppose researchers can find a better set of observations and reward metrics with what already exists in this model. In that case it can only help when the model is expanded or improved with higher or lower fidelity components.

The policy might benefit from adding a time-related observation or reward metric. If the DRL policy is trained to minimize the length of the scenario run, it could produce a better result. Adding a reward metric that subtracts reward the longer the scenario runs might incentivize a different outcome each timestep. The DRL should learn that guiding civilians to the exits quicker results in a higher long-term reward score. This might also be expressed in a reward metric relative to the number of civilians that are still alive in the scenario and the time expired. Therefore, instead of simply having the number of survivors as a metric, it would be the number of survivors divided by the current time. Then, the longer the scenario runs, the lower the score will be. The number of survivors also needs to remain high for a higher reward function result.

Further, adding an observation of the distance each civilian has to the nearest safe exit might help training. However, given that the civilian location and the location of each exit is already present in the current set of observations, this is most likely already captured within the complexity of the neural network. Though, the author has discovered that even the most obvious and well-though-out observations often yield surprising results when dealing with DRL policy training. Therefore, given the experimental nature of this work, it cannot hurt to attempt this and other sets of observations, given enough time. The distance to the closest safe exit might also work as a reward metric. A function of each individual civilian and their distance to the exit, balanced with the other metrics, could add extra value in training.

The experimentation with DRL policy training is nearly endless, though the items mentioned above appear most valuable to consider. Certainly, if a researcher spends more time in thought on this issue, many more different combinations of observations and reward metrics can be deduced and implementing these changes might yield improved results after developing the model further.

5.4.2 Improving the Model

The model itself would benefit from several improvements that ultimately assist in training a better DRL policy and move toward a real-world application. Having a model that simulates a real-world environment as closely as possible to train a DRL policy that controls the dynamic signs is the final goal. Then, the policy can be applied to the true environment within a defensive system. Implied is the addition in observation, reward metrics, or action space needed to reflect the model improvements within DRL policy training.

The first step in this quest should be the implementation of directional control of the dynamic signs. According to the patent listed in previous research, each sign will be able to point an individual in the direction of safety or the next sign to follow (Frantz, 2021). This is a crucial addition to the model since it directly reflects the capabilities of a real-world system. With this addition, changes to the DRL observations and action-space must be made so that the DRL policy can be trained to point each sign in the appropriate direction. Also, civilian agents must be modified to face in the correct direction based on what the signs suggest. This would inherently add the higher level of fidelity to the model, required for real-world application.

The improvement of civilian actions and behavior is the next step. Modeling civilian crowd behavior is important in training the DRL properly and to measure results realistically. A basic model of panic and confusion that can be validated with a historic incident could add the level of fidelity needed to improve the effectiveness of the DRL-controlled signs when applying them to the real world. However, caution must be taken to avoid implementing a system that is too complex and cost prohibitive to develop, missing the mark in regard to real-world applicability. It must have enough fidelity so that the DRL policy can account for realistic human behavior, but not so much that it cannot be trained efficiently or effectively. The concept of abstraction would benefit implementation in this case.

Modeling a higher fidelity visual detection system for civilians is also something to consider. Basing the field of view the civilian agents use on strict scientific literature is an important addition to the model. This level of accuracy was not needed for the current research but is more important for a system that should function in the real world. With visual detection, auditory modeling might be necessary, even if it is abstracted. Given that the real-world signs are able to project sound, this important aspect should be captured in the model. As always, caution must be taken to not over-engineer the model to avoid inefficient and ineffective training of the policy. If the end goal is to deploy a trained policy to a real-world system, the correct balance in the fidelity of the model and complexity in DRL parameters for training must be found. Within this area of improvement, the civilian agent should also include logic that detects obstacles (such as tents, buildings, fencing, etc.) and chooses to avoid looking in the direction of those barriers for a sign to follow. This would add fidelity to the model and yield a DRL policy that better fits a real work environment.

Lastly, being able to train the DRL policy with a random shooter spawn could be of value. Expanding the spawn locations beyond four, based on the cardinal directions, then adding the ability to spawn the shooter randomly, should offer an interesting level of complexity to training the DRL policy while creating a more realistic approach to the model.

5.4.3 A Note on Fidelity and Abstraction

There must be a balance in the level of fidelity of the model and in the complexity of the observation, reward metric and action space of the DRL policy training. Too high a fidelity, and the DRL policy will take excessively long to train or not train effectively. Too low a fidelity, and it might not yield favorable results even if it trained over a long period of time. In addition, too many observations or the wrong action space could produce ineffective results.

Sometimes, fidelity needed in the model can be maintained without taxing processing power by abstracting away complexity. An example related to the model used in this work is in relation to civilian agents. The author chose to model each individual civilian agent. This, however, reduces the ability to model large crowds of potentially thousands since the systems used are limited in how many agents can be modeled concurrently. Even if they were able to use many more individual agents, it could take so long to train a policy that it might become unwise even to attempt it. Hence, abstracting away the details of individual agents could work as a solution. The author suggests using “clusters” of civilians with a certain size or mass instead of individual civilians. These clusters or groups of civilian agents could have their own properties of movement and detection as well as other important aspects. They would form around gathering areas of an open-air venue, and, when the incident starts, move as a cluster. This would also allow the reduction of fidelity in a shooter agent since it would not have to acquire and engage individual agents but only the cluster. A modeler can then abstract the complexity of engagement of each civilian agent while still maintaining the real-world aspect of causing casualties. Overall, these abstractions can reduce the need for processor cycles during the model runs, and most importantly, during DRL policy training. This, by default, can reduce the overall observation space fed to the DRL neural network, given the reduction from potentially thousands of civilian agents to a few-dozen clusters of the same.

APPENDIX A. EXPERIMENT SUMMARY DATA

North Spawn	Mean Runtime	Mean Survivors	Std. Dev Survivors	Mean Casualties	Std. Dev Casualties	Difference In Casualties (%)
Run to Exit	323.18	39.31	3.76	60.69	3.76	-31.92
All Signs Active	424.52	32.34	3.92	67.66	3.92	-38.94
Policy Enabled	390.75	58.69	5.17	41.31	5.17	

South Spawn	Mean Runtime	Mean Survivors	Std. Dev Survivors	Mean Casualties	Std. Dev Casualties	Difference In Casualties (%)
Run to Exit	323.43	84.99	4.21	15.01	4.21	23.35
All Signs Active	378.74	34.87	4.32	65.13	4.32	-71.56
Policy Enabled	367.20	81.48	4.89	18.52	4.89	

East Spawn	Mean Runtime	Mean Survivors	Std. Dev Survivors	Mean Casualties	Std. Dev Casualties	Difference In Casualties (%)
Run to Exit	323.26	60.73	4.62	39.27	4.62	-19.63
All Signs Active	46.65	43.00	5.00	57.00	5.00	-44.62
Policy Enabled	384.44	68.44	4.93	31.56	4.93	

West Spawn	Mean Runtime	Mean Survivors	Std. Dev Survivors	Mean Casualties	Std. Dev Casualties	Difference In Casualties (%)
Run to Exit	327.84	88.89	3.52	11.11	3.52	-91.30
All Signs Active	416.62	71.38	5.06	28.62	5.06	-96.62
Policy Enabled	410.97	99.03	2.04	0.97	2.04	

APPENDIX B. SPAWN RELATED SCENARIO RUN DATA

Iteration	North Spawn		South Spawn		East Spawn		West Spawn	
	Survivors	Casualties	Survivors	Casualties	Survivors	Casualties	Survivors	Casualties
1	55	45	58	42	62	38	96	4
2	71	29	73	27	62	38	96	4
3	61	39	80	20	65	35	100	0
4	63	37	87	13	73	27	100	0
5	56	44	76	24	68	32	100	0
6	58	42	84	16	68	32	100	0
7	49	51	84	16	68	32	98	2
8	65	35	81	19	65	35	94	6
9	60	40	80	20	74	26	100	0
10	56	44	82	18	63	37	95	5
11	60	40	83	17	71	29	99	1
12	65	35	70	30	67	33	100	0
13	47	53	88	12	72	28	100	0
14	58	42	84	16	60	40	100	0
15	62	38	71	29	67	33	100	0
16	61	39	76	24	76	24	99	1
17	64	36	86	14	62	38	100	0
18	61	39	85	15	74	26	100	0
19	57	43	83	17	61	39	100	0
20	62	38	81	19	76	24	100	0
21	60	40	85	15	71	29	100	0
22	66	34	86	14	78	22	100	0
23	62	38	67	33	70	30	93	7
24	68	32	71	29	69	31	99	1
25	56	44	76	24	68	32	100	0
26	64	36	84	16	67	33	100	0
27	59	41	85	15	61	39	100	0
28	57	43	81	19	72	28	100	0
29	60	40	89	11	75	25	89	11
30	65	35	81	19	68	32	100	0

31	58	42	89	11	70	30	100	0
32	58	42	64	36	66	34	100	0
33	67	33	78	22	72	28	100	0
34	65	35	82	18	71	29	100	0
35	48	52	81	19	66	34	96	4
36	62	38	79	21	77	23	100	0
37	58	42	77	23	66	34	96	4
38	57	43	88	12	62	38	100	0
39	60	40	81	19	66	34	99	1
40	62	38	84	16	62	38	100	0
41	64	36	81	19	64	36	100	0
42	66	34	85	15	56	44	100	0
43	56	44	81	19	66	34	100	0
44	71	29	77	23	68	32	100	0
45	54	46	84	16	65	35	100	0
46	55	45	86	14	73	27	100	0
47	59	41	83	17	69	31	100	0
48	58	42	84	16	69	31	100	0
49	58	42	79	21	65	35	94	6
50	65	35	84	16	70	30	99	1
51	68	32	83	17	69	31	100	0
52	63	37	78	22	70	30	100	0
53	55	45	79	21	66	34	100	0
54	50	50	83	17	73	27	100	0
55	63	37	78	22	71	29	100	0
56	60	40	89	11	74	26	99	1
57	63	37	80	20	69	31	97	3
58	59	41	89	11	64	36	100	0
59	58	42	81	19	63	37	100	0
60	57	43	79	21	73	27	100	0
61	50	50	70	30	71	29	100	0
62	54	46	79	21	70	30	98	2
63	57	43	82	18	58	42	100	0
64	54	46	76	24	73	27	96	4
65	58	42	88	12	60	40	100	0
66	63	37	81	19	63	37	100	0
67	64	36	77	23	70	30	100	0

68	53	47	86	14	71	29	100	0
69	62	38	90	10	75	25	99	1
70	55	45	82	18	68	32	100	0
71	53	47	75	25	59	41	100	0
72	47	53	81	19	65	35	99	1
73	62	38	81	19	65	35	100	0
74	63	37	85	15	68	32	94	6
75	63	37	79	21	56	44	100	0
76	59	41	85	15	75	25	99	1
77	60	40	80	20	69	31	100	0
78	48	52	79	21	63	37	98	2
79	56	44	83	17	69	31	100	0
80	59	41	80	20	67	33	96	4
81	66	34	73	27	71	29	100	0
82	68	32	82	18	70	30	100	0
83	62	38	79	21	70	30	100	0
84	55	45	81	19	67	33	100	0
85	62	38	76	24	70	30	100	0
86	65	35	85	15	77	23	100	0
87	57	43	69	31	65	35	100	0
88	56	44	86	14	68	32	91	9
89	63	37	82	18	68	32	100	0
90	46	54	84	16	58	42	100	0
91	54	46	75	25	71	29	100	0
92	56	44	82	18	73	27	100	0
93	61	39	69	31	78	22	100	0
94	65	35	80	20	70	30	100	0
95	64	36	81	19	69	31	95	5
96	64	36	88	12	74	26	100	0
97	58	42	72	28	66	34	100	0
98	54	46	90	10	62	38	100	0
99	56	44	92	8	72	28	100	0
100	57	43	81	19	60	40	100	0
101	60	40	75	25	73	27	100	0
102	53	47	83	17	73	27	97	3
103	59	41	82	18	72	28	100	0
104	60	40	85	15	74	26	99	1

105	57	43	85	15	67	33	100	0
106	60	40	85	15	66	34	100	0
107	52	48	75	25	70	30	99	1
108	58	42	81	19	67	33	99	1
109	67	33	90	10	58	42	100	0
110	55	45	87	13	60	40	100	0
111	53	47	85	15	68	32	100	0
112	56	44	73	27	64	36	100	0
113	58	42	80	20	73	27	99	1
114	59	41	84	16	70	30	100	0
115	60	40	83	17	65	35	100	0
116	64	36	84	16	62	38	100	0
117	55	45	81	19	68	32	99	1
118	58	42	81	19	71	29	100	0
119	41	59	83	17	66	34	100	0
120	67	33	79	21	60	40	100	0
121	57	43	76	24	70	30	100	0
122	64	36	85	15	69	31	96	4
123	63	37	83	17	72	28	93	7
124	53	47	92	8	60	40	100	0
125	60	40	83	17	63	37	99	1
126	60	40	85	15	76	24	100	0
127	60	40	81	19	67	33	100	0
128	66	34	81	19	75	25	100	0
129	58	42	75	25	74	26	98	2
130	66	34	79	21	67	33	100	0
131	58	42	80	20	62	38	100	0
132	60	40	83	17	75	25	100	0
133	61	39	85	15	77	23	100	0
134	66	34	81	19	76	24	98	2
135	64	36	81	19	64	36	100	0
136	61	39	86	14	71	29	100	0
137	59	41	80	20	58	42	100	0
138	61	39	85	15	66	34	100	0
139	55	45	79	21	66	34	100	0
140	60	40	83	17	71	29	100	0
141	61	39	80	20	66	34	94	6

142	64	36	81	19	65	35	90	10
143	69	31	86	14	66	34	100	0
144	58	42	80	20	66	34	98	2
145	58	42	76	24	64	36	100	0
146	59	41	77	23	66	34	100	0
147	60	40	83	17	78	22	100	0
148	63	37	83	17	72	28	100	0
149	60	40	86	14	59	41	100	0
150	62	38	78	22	63	37	96	4
151	59	41	82	18	67	33	100	0
152	60	40	78	22	71	29	96	4
153	59	41	82	18	65	35	100	0
154	58	42	84	16	71	29	96	4
155	62	38	74	26	73	27	89	11
156	49	51	88	12	74	26	100	0
157	60	40	73	27	57	43	100	0
158	56	44	83	17	71	29	100	0
159	62	38	82	18	74	26	100	0
160	57	43	86	14	65	35	100	0
161	61	39	82	18	72	28	89	11
162	58	42	81	19	78	22	100	0
163	57	43	80	20	61	39	100	0
164	61	39	85	15	68	32	98	2
165	63	37	73	27	62	38	100	0
166	59	41	86	14	58	42	99	1
167	60	40	84	16	66	34	100	0
168	48	52	87	13	57	43	100	0
169	65	35	82	18	72	28	99	1
170	58	42	79	21	68	32	100	0
171	53	47	83	17	73	27	100	0
172	58	42	83	17	79	21	100	0
173	64	36	87	13	69	31	100	0
174	71	29	79	21	73	27	100	0
175	56	44	82	18	76	24	92	8
176	62	38	88	12	68	32	100	0
177	64	36	75	25	61	39	100	0
178	60	40	82	18	70	30	100	0

179	54	46	83	17	69	31	100	0
180	58	42	81	19	68	32	100	0
181	62	38	93	7	65	35	100	0
182	61	39	84	16	68	32	100	0
183	66	34	78	22	64	36	100	0
184	64	36	84	16	66	34	100	0
185	66	34	80	20	65	35	100	0
186	59	41	85	15	67	33	100	0
187	62	38	77	23	70	30	97	3
188	57	43	87	13	66	34	100	0
189	64	36	84	16	79	21	100	0
190	51	49	89	11	70	30	100	0
191	58	42	82	18	64	36	100	0
192	58	42	82	18	76	24	100	0
193	53	47	84	16	75	25	100	0
194	58	42	81	19	71	29	100	0
195	63	37	72	28	69	31	100	0
196	59	41	76	24	71	29	100	0
197	58	42	85	15	59	41	100	0
198	63	37	87	13	51	49	100	0
199	65	35	77	23	72	28	97	3
200	61	39	83	17	66	34	99	1
201	56	44	88	12	62	38	100	0
202	56	44	86	14	64	36	99	1
203	64	36	91	9	69	31	100	0
204	49	51	83	17	70	30	100	0
205	65	35	86	14	56	44	99	1
206	66	34	74	26	70	30	94	6
207	57	43	85	15	67	33	99	1
208	58	42	83	17	65	35	99	1
209	52	48	84	16	68	32	100	0
210	64	36	70	30	73	27	100	0
211	53	47	81	19	65	35	99	1
212	55	45	75	25	68	32	100	0
213	64	36	86	14	68	32	100	0
214	61	39	84	16	75	25	100	0
215	61	39	83	17	70	30	100	0

216	60	40	87	13	63	37	100	0
217	61	39	79	21	56	44	100	0
218	73	27	81	19	60	40	94	6
219	55	45	77	23	76	24	100	0
220	46	54	82	18	68	32	99	1
221	63	37	88	12	52	48	100	0
222	55	45	87	13	70	30	100	0
223	51	49	76	24	75	25	100	0
224	62	38	71	29	64	36	90	10
225	57	43	76	24	71	29	99	1
226	57	43	88	12	69	31	100	0
227	56	44	86	14	75	25	100	0
228	51	49	74	26	70	30	100	0
229	46	54	85	15	59	41	94	6
230	61	39	87	13	56	44	100	0
231	56	44	86	14	79	21	100	0
232	59	41	82	18	66	34	100	0
233	62	38	72	28	69	31	100	0
234	62	38	83	17	72	28	100	0
235	63	37	88	12	70	30	96	4
236	45	55	80	20	63	37	94	6
237	62	38	85	15	71	29	100	0
238	59	41	83	17	70	30	100	0
239	64	36	84	16	68	32	99	1
240	59	41	82	18	71	29	100	0
241	54	46	80	20	80	20	98	2
242	62	38	85	15	66	34	100	0
243	61	39	83	17	65	35	100	0
244	60	40	83	17	65	35	100	0
245	51	49	84	16	55	45	100	0
246	52	48	81	19	72	28	100	0
247	57	43	87	13	72	28	100	0
248	63	37	81	19	69	31	100	0
249	56	44	74	26	64	36	100	0
250	62	38	83	17	70	30	100	0
251	61	39	77	23	70	30	94	6
252	58	42	82	18	69	31	100	0

253	62	38	64	36	68	32	100	0
254	61	39	89	11	65	35	100	0
255	62	38	86	14	62	38	96	4
256	57	43	83	17	64	36	100	0
257	60	40	82	18	72	28	97	3
258	61	39	79	21	73	27	100	0
259	60	40	82	18	63	37	93	7
260	64	36	60	40	69	31	98	2
261	57	43	88	12	70	30	100	0
262	64	36	75	25	68	32	100	0
263	60	40	82	18	75	25	100	0
264	56	44	84	16	73	27	100	0
265	59	41	83	17	66	34	100	0
266	65	35	83	17	72	28	98	2
267	48	52	88	12	65	35	100	0
268	64	36	89	11	68	32	100	0
269	60	40	89	11	66	34	98	2
270	64	36	74	26	75	25	100	0
271	55	45	83	17	78	22	100	0
272	64	36	82	18	71	29	98	2
273	64	36	86	14	72	28	98	2
274	54	46	83	17	72	28	100	0
275	56	44	84	16	68	32	96	4
276	51	49	78	22	67	33	100	0
277	61	39	78	22	65	35	100	0
278	59	41	83	17	62	38	97	3
279	63	37	92	8	65	35	100	0
280	64	36	84	16	71	29	100	0
281	59	41	81	19	65	35	100	0
282	56	44	81	19	62	38	100	0
283	65	35	87	13	68	32	100	0
284	62	38	82	18	64	36	96	4
285	61	39	82	18	70	30	100	0
286	52	48	83	17	74	26	100	0
287	39	61	84	16	71	29	98	2
288	62	38	74	26	66	34	100	0
289	64	36	84	16	62	38	98	2

290	67	33	72	28	70	30	100	0
291	55	45	84	16	69	31	98	2
292	64	36	81	19	77	23	100	0
293	59	41	78	22	82	18	99	1
294	55	45	73	27	61	39	100	0
295	59	41	92	8	69	31	97	3
296	54	46	82	18	59	41	99	1
297	58	42	81	19	62	38	100	0
298	56	44	82	18	67	33	97	3
299	62	38	85	15	66	34	99	1
300	52	48	78	22	70	30	95	5
301	60	40	81	19	66	34	100	0
302	47	53	87	13	71	29	100	0
303	58	42	87	13	72	28	99	1
304	59	41	73	27	75	25	93	7
305	57	43	82	18	72	28	98	2
306	61	39	82	18	70	30	100	0
307	65	35	86	14	64	36	100	0
308	48	52	87	13	73	27	100	0
309	53	47	83	17	72	28	100	0
310	60	40	73	27	68	32	97	3
311	56	44	88	12	73	27	91	9
312	64	36	85	15	72	28	100	0
313	61	39	82	18	61	39	100	0
314	57	43	77	23	75	25	100	0
315	59	41	84	16	74	26	100	0
316	61	39	80	20	76	24	100	0
317	59	41	82	18	73	27	100	0
318	58	42	77	23	64	36	100	0
319	66	34	86	14	75	25	99	1
320	64	36	86	14	70	30	100	0
321	65	35	88	12	68	32	100	0
322	63	37	84	16	81	19	100	0
323	71	29	83	17	74	26	100	0
324	51	49	84	16	68	32	96	4
325	59	41	89	11	62	38	97	3
326	56	44	83	17	70	30	88	12

327	59	41	81	19	70	30	100	0
328	56	44	70	30	70	30	97	3
329	59	41	80	20	64	36	96	4
330	69	31	82	18	71	29	100	0
331	60	40	81	19	70	30	99	1
332	49	51	80	20	70	30	100	0
333	60	40	87	13	64	36	100	0
334	57	43	82	18	76	24	100	0
335	53	47	81	19	64	36	100	0
336	62	38	85	15	73	27	93	7
337	64	36	81	19	66	34	92	8
338	58	42	85	15	69	31	100	0
339	57	43	81	19	62	38	100	0
340	59	41	84	16	70	30	100	0
341	59	41	91	9	66	34	97	3
342	56	44	86	14	70	30	100	0
343	57	43	77	23	62	38	100	0
344	59	41	87	13	62	38	100	0
345	53	47	81	19	67	33	100	0
346	60	40	75	25	70	30	100	0
347	61	39	81	19	77	23	100	0
348	55	45	74	26	74	26	99	1
349	60	40	81	19	68	32	99	1
350	60	40	83	17	66	34	100	0
351	54	46	87	13	63	37	93	7
352	58	42	87	13	66	34	100	0
353	48	52	77	23	76	24	100	0
354	55	45	89	11	73	27	100	0
355	67	33	80	20	68	32	99	1
356	56	44	74	26	77	23	100	0
357	57	43	85	15	74	26	100	0
358	54	46	76	24	61	39	100	0
359	65	35	84	16	63	37	100	0
360	61	39	85	15	60	40	98	2
361	62	38	81	19	69	31	95	5
362	59	41	84	16	71	29	96	4
363	46	54	80	20	61	39	99	1

364	65	35	82	18	73	27	96	4
365	57	43	82	18	68	32	98	2
366	61	39	79	21	71	29	100	0
367	59	41	85	15	71	29	100	0
368	59	41	79	21	67	33	100	0
369	41	59	80	20	71	29	100	0
370	47	53	87	13	63	37	100	0
371	59	41	81	19	65	35	100	0
372	57	43	90	10	69	31	100	0
373	64	36	82	18	66	34	92	8
374	41	59	83	17	70	30	100	0
375	53	47	85	15	62	38	96	4
376	62	38	86	14	71	29	100	0
377	62	38	83	17	72	28	100	0
378	55	45	92	8	74	26	100	0
379	53	47	74	26	62	38	100	0
380	58	42	86	14	64	36	100	0
381	63	37	86	14	66	34	100	0
382	59	41	84	16	73	27	100	0
383	64	36	80	20	68	32	100	0
384	59	41	81	19	68	32	100	0
385	63	37	85	15	67	33	100	0
386	66	34	79	21	78	22	100	0
387	59	41	88	12	75	25	100	0
388	57	43	76	24	67	33	100	0
389	52	48	70	30	75	25	95	5
390	56	44	89	11	62	38	95	5
391	62	38	88	12	63	37	100	0
392	60	40	93	7	61	39	100	0
393	64	36	81	19	63	37	100	0
394	56	44	85	15	63	37	99	1
395	47	53	80	20	63	37	98	2
396	54	46	79	21	79	21	97	3
397	60	40	83	17	68	32	100	0
398	51	49	87	13	69	31	100	0
399	61	39	80	20	63	37	92	8
400	66	34	78	22	74	26	98	2

401	59	41	85	15	72	28	99	1
402	55	45	74	26	71	29	100	0
403	63	37	84	16	62	38	100	0
404	58	42	76	24	66	34	100	0
405	58	42	81	19	67	33	100	0
406	51	49	80	20	75	25	100	0
407	50	50	82	18	79	21	100	0
408	55	45	81	19	60	40	99	1
409	64	36	81	19	64	36	100	0
410	61	39	85	15	66	34	90	10
411	61	39	84	16	65	35	98	2
412	60	40	83	17	70	30	100	0
413	63	37	90	10	71	29	100	0
414	56	44	66	34	68	32	100	0
415	64	36	72	28	68	32	100	0
416	58	42	78	22	62	38	100	0
417	60	40	80	20	76	24	96	4
418	66	34	80	20	67	33	100	0
419	56	44	86	14	69	31	100	0
420	57	43	82	18	66	34	100	0
421	59	41	81	19	67	33	99	1
422	62	38	80	20	70	30	100	0
423	55	45	88	12	69	31	100	0
424	53	47	81	19	66	34	100	0
425	55	45	81	19	69	31	100	0
426	61	39	83	17	74	26	100	0
427	58	42	78	22	69	31	100	0
428	58	42	79	21	72	28	100	0
429	59	41	74	26	63	37	100	0
430	58	42	83	17	73	27	100	0
431	61	39	89	11	64	36	100	0
432	62	38	86	14	68	32	100	0
433	60	40	81	19	57	43	99	1
434	57	43	86	14	68	32	100	0
435	49	51	79	21	71	29	100	0
436	61	39	83	17	69	31	100	0
437	56	44	71	29	70	30	100	0

438	62	38	77	23	72	28	100	0
439	60	40	77	23	69	31	98	2
440	70	30	76	24	68	32	100	0
441	35	65	85	15	68	32	100	0
442	62	38	81	19	67	33	100	0
443	60	40	77	23	68	32	100	0
444	58	42	93	7	62	38	100	0
445	62	38	90	10	67	33	100	0
446	63	37	78	22	65	35	100	0
447	60	40	84	16	72	28	100	0
448	52	48	83	17	69	31	100	0
449	51	49	72	28	74	26	100	0
450	56	44	80	20	65	35	100	0
451	58	42	86	14	72	28	100	0
452	65	35	88	12	59	41	96	4
453	65	35	82	18	71	29	94	6
454	53	47	89	11	66	34	100	0
455	52	48	79	21	68	32	100	0
456	60	40	85	15	75	25	100	0
457	54	46	84	16	68	32	100	0
458	58	42	82	18	77	23	100	0
459	55	45	67	33	71	29	100	0
460	57	43	79	21	65	35	100	0
461	59	41	81	19	63	37	100	0
462	62	38	65	35	59	41	100	0
463	50	50	90	10	67	33	100	0
464	41	59	85	15	72	28	95	5
465	58	42	80	20	60	40	100	0
466	60	40	84	16	69	31	98	2
467	66	34	87	13	75	25	99	1
468	60	40	75	25	67	33	100	0
469	61	39	84	16	65	35	100	0
470	56	44	87	13	58	42	100	0
471	62	38	84	16	67	33	100	0
472	61	39	80	20	67	33	90	10
473	53	47	82	18	69	31	100	0
474	57	43	86	14	73	27	99	1

475	57	43	78	22	64	36	98	2
476	55	45	79	21	71	29	99	1
477	58	42	78	22	73	27	98	2
478	63	37	83	17	61	39	100	0
479	54	46	83	17	72	28	100	0
480	60	40	76	24	69	31	95	5
481	52	48	80	20	69	31	100	0
482	55	45	81	19	64	36	100	0
483	58	42	84	16	72	28	99	1
484	56	44	75	25	79	21	99	1
485	59	41	77	23	69	31	99	1
486	51	49	79	21	64	36	100	0
487	68	32	76	24	77	23	96	4
488	56	44	73	27	65	35	100	0
489	57	43	74	26	73	27	100	0
490	58	42	83	17	66	34	94	6
491	59	41	90	10	72	28	100	0
492	56	44	78	22	56	44	100	0
493	65	35	78	22	76	24	99	1
494	57	43	80	20	62	38	100	0
495	45	55	87	13	70	30	100	0
496	61	39	79	21	63	37	100	0
497	57	43	88	12	72	28	97	3
498	56	44	85	15	63	37	98	2
499	52	48	87	13	70	30	100	0
500	54	46	82	18	72	28	94	6
501	58	42	79	21	72	28	100	0
502	55	45	88	12	70	30	100	0
503	63	37	79	21	65	35	100	0
504	62	38	82	18	74	26	100	0
505	63	37	74	26	64	36	100	0
506	56	44	86	14	65	35	100	0
507	62	38	88	12	76	24	100	0
508	57	43	77	23	68	32	100	0
509	63	37	82	18	60	40	100	0
510	55	45	87	13	66	34	100	0
511	53	47	82	18	69	31	98	2

512	54	46	83	17	70	30	98	2
513	58	42	69	31	66	34	100	0
514	62	38	81	19	71	29	100	0
515	54	46	81	19	66	34	100	0
516	61	39	83	17	66	34	100	0
517	66	34	78	22	71	29	98	2
518	62	38	82	18	75	25	100	0
519	58	42	80	20	67	33	97	3
520	42	58	80	20	65	35	100	0
521	57	43	81	19	71	29	100	0
522	68	32	81	19	73	27	96	4
523	61	39	88	12	68	32	100	0
524	63	37	79	21	70	30	100	0
525	66	34	74	26	59	41	100	0
526	59	41	79	21	66	34	100	0
527	55	45	81	19	69	31	100	0
528	48	52	73	27	70	30	99	1
529	62	38	78	22	59	41	99	1
530	55	45	78	22	73	27	100	0
531	59	41	81	19	69	31	100	0
532	64	36	86	14	72	28	100	0
533	54	46	84	16	76	24	98	2
534	64	36	74	26	65	35	100	0
535	51	49	83	17	70	30	99	1
536	56	44	80	20	66	34	100	0
537	60	40	80	20	62	38	100	0
538	61	39	78	22	68	32	98	2
539	61	39	82	18	65	35	99	1
540	55	45	79	21	65	35	95	5
541	60	40	87	13	64	36	100	0
542	60	40	67	33	65	35	100	0
543	68	32	77	23	64	36	99	1
544	61	39	94	6	69	31	100	0
545	63	37	83	17	75	25	100	0
546	54	46	83	17	69	31	98	2
547	60	40	77	23	73	27	100	0
548	58	42	83	17	61	39	100	0

549	56	44	89	11	70	30	100	0
550	65	35	86	14	69	31	100	0
551	56	44	78	22	67	33	99	1
552	63	37	74	26	69	31	100	0
553	59	41	81	19	60	40	97	3
554	64	36	79	21	74	26	100	0
555	64	36	78	22	70	30	98	2
556	70	30	84	16	44	56	100	0
557	59	41	78	22	70	30	99	1
558	53	47	83	17	73	27	100	0
559	47	53	84	16	70	30	90	10
560	62	38	79	21	67	33	100	0
561	63	37	81	19	70	30	100	0
562	58	42	83	17	78	22	100	0
563	60	40	85	15	73	27	100	0
564	62	38	81	19	67	33	100	0
565	63	37	84	16	66	34	100	0
566	59	41	79	21	64	36	100	0
567	64	36	85	15	75	25	100	0
568	53	47	83	17	64	36	100	0
569	58	42	69	31	67	33	100	0
570	54	46	76	24	73	27	99	1
571	60	40	82	18	65	35	100	0
572	60	40	76	24	66	34	100	0
573	64	36	68	32	69	31	92	8
574	53	47	83	17	71	29	100	0
575	54	46	89	11	67	33	100	0
576	57	43	77	23	63	37	100	0
577	54	46	80	20	66	34	100	0
578	63	37	82	18	71	29	100	0
579	59	41	80	20	67	33	100	0
580	62	38	89	11	63	37	100	0
581	57	43	81	19	68	32	100	0
582	63	37	83	17	70	30	100	0
583	65	35	86	14	66	34	100	0
584	57	43	86	14	68	32	95	5
585	65	35	75	25	69	31	99	1

586	61	39	80	20	69	31	100	0
587	57	43	81	19	73	27	100	0
588	62	38	80	20	79	21	95	5
589	56	44	80	20	74	26	100	0
590	66	34	83	17	65	35	100	0
591	64	36	80	20	70	30	99	1
592	67	33	77	23	70	30	100	0
593	61	39	85	15	61	39	96	4
594	64	36	76	24	69	31	100	0
595	61	39	83	17	79	21	98	2
596	56	44	81	19	60	40	96	4
597	57	43	83	17	70	30	96	4
598	54	46	87	13	60	40	100	0
599	63	37	71	29	77	23	94	6
600	64	36	80	20	75	25	100	0
601	56	44	81	19	66	34	98	2
602	64	36	87	13	59	41	100	0
603	60	40	82	18	72	28	100	0
604	57	43	88	12	62	38	92	8
605	56	44	82	18	74	26	94	6
606	47	53	83	17	64	36	100	0
607	51	49	71	29	63	37	98	2
608	60	40	84	16	69	31	100	0
609	60	40	74	26	75	25	90	10
610	61	39	74	26	73	27	100	0
611	67	33	73	27	69	31	100	0
612	64	36	80	20	59	41	93	7
613	70	30	76	24	75	25	100	0
614	60	40	81	19	74	26	100	0
615	63	37	89	11	66	34	100	0
616	60	40	68	32	67	33	100	0
617	61	39	77	23	69	31	100	0
618	60	40	87	13	70	30	99	1
619	62	38	78	22	72	28	100	0
620	65	35	79	21	71	29	99	1
621	59	41	82	18	74	26	100	0
622	52	48	87	13	67	33	100	0

623	50	50	90	10	68	32	99	1
624	59	41	74	26	58	42	100	0
625	65	35	81	19	65	35	96	4
626	58	42	83	17	75	25	100	0
627	60	40	78	22	66	34	93	7
628	55	45	85	15	74	26	100	0
629	52	48	82	18	70	30	98	2
630	59	41	74	26	59	41	100	0
631	61	39	86	14	78	22	96	4
632	64	36	79	21	70	30	100	0
633	64	36	84	16	65	35	93	7
634	65	35	82	18	74	26	100	0
635	57	43	86	14	58	42	96	4
636	59	41	85	15	62	38	99	1
637	57	43	75	25	65	35	100	0
638	60	40	85	15	70	30	100	0
639	58	42	79	21	71	29	100	0
640	48	52	79	21	67	33	100	0
641	61	39	83	17	77	23	100	0
642	67	33	82	18	71	29	100	0
643	53	47	84	16	72	28	100	0
644	55	45	80	20	63	37	100	0
645	56	44	82	18	70	30	100	0
646	59	41	87	13	68	32	98	2
647	64	36	80	20	66	34	100	0
648	47	53	72	28	74	26	100	0
649	52	48	75	25	67	33	100	0
650	53	47	81	19	70	30	100	0
651	56	44	78	22	77	23	100	0
652	55	45	89	11	63	37	94	6
653	60	40	86	14	69	31	100	0
654	58	42	75	25	77	23	95	5
655	59	41	80	20	67	33	97	3
656	56	44	91	9	65	35	99	1
657	57	43	83	17	69	31	100	0
658	64	36	87	13	69	31	100	0
659	51	49	82	18	66	34	100	0

660	63	37	86	14	68	32	100	0
661	57	43	80	20	71	29	100	0
662	64	36	77	23	68	32	100	0
663	53	47	84	16	65	35	100	0
664	55	45	78	22	71	29	100	0
665	63	37	83	17	60	40	91	9
666	51	49	84	16	63	37	99	1
667	60	40	85	15	63	37	99	1
668	54	46	79	21	71	29	100	0
669	57	43	78	22	71	29	100	0
670	59	41	73	27	68	32	100	0
671	61	39	84	16	62	38	100	0
672	55	45	80	20	66	34	99	1
673	65	35	80	20	71	29	93	7
674	62	38	90	10	66	34	95	5
675	55	45	80	20	68	32	99	1
676	46	54	81	19	58	42	95	5
677	59	41	82	18	61	39	100	0
678	50	50	81	19	73	27	100	0
679	55	45	82	18	70	30	100	0
680	55	45	74	26	67	33	100	0
681	59	41	86	14	67	33	100	0
682	62	38	76	24	67	33	99	1
683	59	41	84	16	63	37	95	5
684	62	38	80	20	66	34	100	0
685	61	39	78	22	71	29	100	0
686	63	37	91	9	67	33	98	2
687	57	43	82	18	57	43	100	0
688	63	37	84	16	68	32	100	0
689	58	42	83	17	75	25	100	0
690	68	32	85	15	72	28	92	8
691	60	40	76	24	74	26	100	0
692	64	36	84	16	74	26	100	0
693	59	41	85	15	70	30	100	0
694	52	48	82	18	70	30	99	1
695	60	40	84	16	74	26	100	0
696	60	40	83	17	71	29	100	0

697	59	41	86	14	72	28	100	0
698	54	46	70	30	70	30	100	0
699	64	36	86	14	69	31	100	0
700	59	41	78	22	72	28	100	0
701	54	46	70	30	66	34	100	0
702	59	41	86	14	68	32	100	0
703	54	46	76	24	76	24	97	3
704	60	40	75	25	79	21	96	4
705	55	45	75	25	72	28	98	2
706	61	39	84	16	72	28	100	0
707	57	43	80	20	68	32	100	0
708	44	56	80	20	67	33	99	1
709	58	42	83	17	64	36	99	1
710	57	43	84	16	72	28	100	0
711	58	42	85	15	64	36	100	0
712	44	56	82	18	72	28	100	0
713	63	37	84	16	68	32	100	0
714	58	42	90	10	68	32	100	0
715	56	44	77	23	65	35	98	2
716	59	41	80	20	70	30	100	0
717	61	39	83	17	76	24	100	0
718	59	41	86	14	71	29	99	1
719	56	44	83	17	68	32	100	0
720	61	39	81	19	72	28	100	0
721	64	36	92	8	70	30	96	4
722	47	53	80	20	63	37	100	0
723	61	39	73	27	70	30	100	0
724	60	40	85	15	71	29	100	0
725	60	40	83	17	69	31	100	0
726	52	48	86	14	75	25	100	0
727	59	41	85	15	74	26	100	0
728	57	43	77	23	69	31	100	0
729	64	36	84	16	71	29	100	0
730	61	39	78	22	71	29	97	3
731	59	41	79	21	71	29	100	0
732	64	36	76	24	71	29	100	0
733	52	48	73	27	71	29	100	0

734	58	42	79	21	80	20	100	0
735	63	37	80	20	73	27	100	0
736	63	37	82	18	75	25	97	3
737	61	39	84	16	68	32	100	0
738	59	41	86	14	71	29	100	0
739	67	33	89	11	66	34	98	2
740	61	39	86	14	75	25	100	0
741	60	40	71	29	73	27	100	0
742	54	46	80	20	67	33	100	0
743	66	34	83	17	72	28	99	1
744	64	36	72	28	64	36	100	0
745	56	44	81	19	77	23	99	1
746	57	43	79	21	64	36	100	0
747	60	40	81	19	61	39	96	4
748	58	42	78	22	65	35	100	0
749	60	40	82	18	76	24	100	0
750	63	37	77	23	74	26	99	1
751	47	53	83	17	68	32	99	1
752	63	37	84	16	69	31	100	0
753	52	48	84	16	71	29	99	1
754	64	36	78	22	67	33	100	0
755	56	44	78	22	68	32	100	0
756	63	37	76	24	52	48	100	0
757	58	42	72	28	70	30	100	0
758	68	32	82	18	71	29	100	0
759	64	36	77	23	68	32	100	0
760	68	32	83	17	56	44	100	0
761	56	44	85	15	70	30	100	0
762	54	46	75	25	63	37	99	1
763	61	39	83	17	67	33	96	4
764	65	35	78	22	70	30	100	0
765	59	41	82	18	75	25	100	0
766	49	51	75	25	62	38	100	0
767	59	41	84	16	67	33	100	0
768	59	41	78	22	72	28	95	5
769	54	46	83	17	63	37	100	0
770	61	39	80	20	66	34	100	0

771	62	38	86	14	69	31	100	0
772	63	37	81	19	65	35	100	0
773	41	59	82	18	63	37	100	0
774	59	41	82	18	63	37	100	0
775	60	40	85	15	70	30	100	0
776	59	41	84	16	71	29	99	1
777	63	37	74	26	73	27	99	1
778	66	34	82	18	71	29	100	0
779	62	38	81	19	68	32	100	0
780	59	41	88	12	70	30	100	0
781	59	41	73	27	64	36	100	0
782	56	44	86	14	68	32	100	0
783	63	37	80	20	70	30	100	0
784	60	40	84	16	57	43	92	8
785	58	42	80	20	74	26	100	0
786	60	40	85	15	73	27	100	0
787	60	40	81	19	71	29	100	0
788	53	47	75	25	63	37	100	0
789	64	36	75	25	60	40	97	3
790	59	41	76	24	63	37	94	6
791	55	45	82	18	69	31	100	0
792	62	38	83	17	73	27	97	3
793	58	42	84	16	68	32	100	0
794	56	44	83	17	70	30	98	2
795	61	39	84	16	60	40	100	0
796	56	44	89	11	71	29	97	3
797	59	41	69	31	64	36	100	0
798	41	59	87	13	66	34	94	6
799	62	38	82	18	66	34	100	0
800	58	42	72	28	66	34	100	0
801	61	39	81	19	59	41	100	0
802	60	40	76	24	65	35	100	0
803	57	43	83	17	77	23	98	2
804	59	41	79	21	70	30	100	0
805	63	37	84	16	63	37	100	0
806	52	48	82	18	72	28	100	0
807	55	45	76	24	70	30	100	0

808	45	55	88	12	73	27	100	0
809	58	42	84	16	75	25	99	1
810	54	46	82	18	64	36	99	1
811	56	44	81	19	72	28	100	0
812	53	47	69	31	65	35	100	0
813	58	42	83	17	69	31	100	0
814	61	39	84	16	69	31	100	0
815	54	46	86	14	66	34	100	0
816	61	39	81	19	73	27	100	0
817	57	43	76	24	66	34	100	0
818	58	42	72	28	76	24	100	0
819	60	40	85	15	66	34	100	0
820	48	52	79	21	69	31	100	0
821	56	44	86	14	70	30	100	0
822	62	38	85	15	73	27	96	4
823	61	39	83	17	68	32	99	1
824	55	45	87	13	66	34	100	0
825	62	38	82	18	63	37	100	0
826	53	47	84	16	67	33	100	0
827	64	36	88	12	76	24	99	1
828	62	38	83	17	68	32	100	0
829	63	37	83	17	74	26	100	0
830	56	44	82	18	69	31	97	3
831	58	42	76	24	71	29	100	0
832	63	37	80	20	74	26	98	2
833	54	46	85	15	78	22	99	1
834	52	48	90	10	65	35	100	0
835	65	35	85	15	68	32	100	0
836	61	39	74	26	66	34	100	0
837	54	46	82	18	70	30	100	0
838	51	49	83	17	73	27	100	0
839	55	45	88	12	72	28	99	1
840	58	42	79	21	73	27	93	7
841	56	44	83	17	67	33	100	0
842	45	55	87	13	72	28	100	0
843	66	34	88	12	69	31	100	0
844	63	37	86	14	66	34	100	0

845	58	42	82	18	70	30	99	1
846	63	37	78	22	66	34	100	0
847	65	35	80	20	66	34	100	0
848	58	42	85	15	74	26	100	0
849	59	41	91	9	74	26	100	0
850	59	41	88	12	69	31	100	0
851	59	41	81	19	67	33	98	2
852	58	42	83	17	73	27	98	2
853	60	40	83	17	68	32	100	0
854	62	38	80	20	75	25	100	0
855	57	43	74	26	65	35	100	0
856	68	32	90	10	67	33	100	0
857	57	43	76	24	62	38	99	1
858	59	41	89	11	66	34	100	0
859	47	53	82	18	66	34	99	1
860	59	41	62	38	66	34	100	0
861	60	40	84	16	78	22	100	0
862	59	41	74	26	73	27	100	0
863	61	39	80	20	61	39	100	0
864	59	41	80	20	70	30	100	0
865	65	35	84	16	71	29	100	0
866	61	39	76	24	56	44	100	0
867	59	41	85	15	71	29	100	0
868	54	46	81	19	67	33	95	5
869	57	43	77	23	66	34	100	0
870	66	34	92	8	70	30	100	0
871	62	38	80	20	68	32	98	2
872	55	45	88	12	79	21	96	4
873	68	32	85	15	72	28	100	0
874	56	44	88	12	70	30	100	0
875	57	43	81	19	73	27	100	0
876	67	33	88	12	68	32	100	0
877	57	43	82	18	71	29	100	0
878	53	47	83	17	64	36	100	0
879	55	45	86	14	75	25	99	1
880	60	40	82	18	66	34	100	0
881	66	34	74	26	70	30	95	5

882	64	36	83	17	68	32	99	1
883	66	34	86	14	68	32	95	5
884	56	44	84	16	65	35	100	0
885	53	47	78	22	75	25	100	0
886	59	41	88	12	61	39	100	0
887	66	34	82	18	76	24	95	5
888	59	41	78	22	63	37	100	0
889	55	45	87	13	65	35	100	0
890	50	50	87	13	74	26	90	10
891	65	35	81	19	65	35	100	0
892	65	35	79	21	75	25	100	0
893	47	53	90	10	71	29	100	0
894	57	43	79	21	68	32	99	1
895	61	39	80	20	78	22	100	0
896	55	45	86	14	68	32	100	0
897	57	43	66	34	76	24	100	0
898	62	38	83	17	82	18	100	0
899	50	50	83	17	70	30	100	0
900	58	42	76	24	73	27	100	0
901	65	35	78	22	65	35	100	0
902	44	56	79	21	66	34	100	0
903	58	42	81	19	68	32	96	4
904	43	57	87	13	69	31	98	2
905	72	28	87	13	71	29	99	1
906	60	40	81	19	73	27	100	0
907	47	53	76	24	71	29	100	0
908	63	37	85	15	72	28	100	0
909	50	50	83	17	65	35	100	0
910	54	46	77	23	63	37	99	1
911	62	38	81	19	63	37	98	2
912	50	50	84	16	72	28	100	0
913	62	38	80	20	67	33	100	0
914	69	31	81	19	75	25	100	0
915	56	44	81	19	71	29	100	0
916	65	35	80	20	70	30	100	0
917	57	43	81	19	73	27	100	0
918	46	54	87	13	81	19	100	0

919	64	36	85	15	61	39	100	0
920	58	42	87	13	69	31	100	0
921	51	49	92	8	63	37	100	0
922	65	35	80	20	70	30	98	2
923	65	35	90	10	69	31	100	0
924	57	43	90	10	66	34	96	4
925	61	39	82	18	68	32	100	0
926	59	41	73	27	70	30	100	0
927	61	39	83	17	69	31	97	3
928	57	43	86	14	70	30	98	2
929	62	38	83	17	63	37	100	0
930	70	30	70	30	61	39	100	0
931	54	46	88	12	74	26	100	0
932	53	47	82	18	69	31	100	0
933	65	35	79	21	71	29	95	5
934	64	36	80	20	64	36	100	0
935	62	38	83	17	71	29	100	0
936	60	40	79	21	60	40	100	0
937	65	35	81	19	75	25	98	2
938	62	38	82	18	64	36	100	0
939	42	58	75	25	61	39	100	0
940	60	40	85	15	67	33	100	0
941	57	43	82	18	70	30	100	0
942	57	43	81	19	69	31	100	0
943	65	35	85	15	69	31	100	0
944	61	39	77	23	65	35	100	0
945	57	43	90	10	76	24	100	0
946	58	42	76	24	63	37	100	0
947	67	33	80	20	72	28	98	2
948	50	50	86	14	72	28	99	1
949	66	34	84	16	68	32	100	0
950	63	37	84	16	64	36	99	1
951	68	32	81	19	72	28	100	0
952	70	30	78	22	73	27	96	4
953	55	45	83	17	67	33	99	1
954	56	44	84	16	74	26	100	0
955	60	40	78	22	66	34	100	0

956	58	42	86	14	71	29	100	0
957	55	45	83	17	68	32	100	0
958	57	43	86	14	69	31	88	12
959	53	47	83	17	69	31	98	2
960	60	40	85	15	66	34	100	0
961	52	48	78	22	78	22	97	3
962	59	41	84	16	63	37	100	0
963	51	49	77	23	67	33	100	0
964	60	40	78	22	75	25	100	0
965	58	42	89	11	65	35	100	0
966	62	38	80	20	71	29	100	0
967	62	38	81	19	65	35	100	0
968	52	48	77	23	68	32	100	0
969	64	36	88	12	74	26	96	4
970	54	46	81	19	72	28	100	0
971	51	49	76	24	65	35	100	0
972	55	45	83	17	69	31	100	0
973	65	35	77	23	70	30	100	0
974	62	38	80	20	67	33	100	0
975	62	38	78	22	60	40	100	0
976	58	42	84	16	72	28	99	1
977	63	37	82	18	67	33	100	0
978	52	48	80	20	66	34	100	0
979	59	41	80	20	55	45	97	3
980	56	44	83	17	75	25	100	0
981	53	47	87	13	72	28	99	1
982	60	40	83	17	70	30	99	1
983	63	37	83	17	69	31	100	0
984	60	40	75	25	66	34	100	0
985	57	43	82	18	69	31	100	0
986	64	36	80	20	73	27	100	0
987	52	48	82	18	67	33	100	0
988	56	44	79	21	66	34	100	0
989	53	47	80	20	76	24	100	0
990	61	39	75	25	76	24	99	1
991	61	39	83	17	73	27	100	0
992	64	36	78	22	69	31	100	0

993	66	34	82	18	69	31	100	0
994	47	53	84	16	68	32	94	6
995	62	38	80	20	71	29	100	0
996	57	43	80	20	70	30	99	1
997	59	41	85	15	71	29	100	0
998	56	44	83	17	59	41	100	0
999	63	37	75	25	69	31	100	0
1000	56	44	81	19	66	34	92	8

APPENDIX C. COMPLETE SOURCE DATA SETS AND MODEL

Contact the author at rbott@purdue.edu for the full set of source data and model components.

REFERENCES

Abar, S., Theodoropoulos, G. K., Lemarinier, P., & O'Hare, G. M. P. (2017). Agent Based Modelling and Simulation tools: A review of the state-of-art software. *Computer Science Review*, *24*, 13–33. https://doi.org/10.1016/j.cosrev.2017.03.001

Alanez, T., Fleshler, D., Hobbs, S., Huriash, L. J., McMahon, P., O'Matz, M., & Travis, S. (2018, December 28). UNPREPARED AND OVERWHELMED. *Www.Sun-Sentinel.Com*. https://projects.sun-sentinel.com/2018/sfl-parkland-school-shooting-critical-moments

ALICE Training Solutions. (n.d.). *ALICE*. Retrieved October 26, 2020, from https://www.alicetraining.com/our-program/

Anklam III, C., Kirby, A., Sharevski, F., & Dietz, J. E. (2015). Mitigating active shooter impact: Analysis for policy options based on agent/computer-based modeling. *Journal of Emergency Management*, *13*(3), 201. https://doi.org/10.5055/jem.2015.0234

AnyLogic. (n.d.). *AnyLogic: Simulation Modeling Software Tools & Solutions for Business*. The AnyLogic Company. Retrieved March 1, 2021, from https://www.anylogic.com

Associated Press. (2016, July 9). Timeline: Dallas Sniper Attack. *Associated Press*. https://spectrumlocalnews.com/tx/austin/news/2016/07/8/timeline--dallas-sniper-attack

ATF. (n.d.). *Bump Stocks | Bureau of Alcohol, Tobacco, Firearms and Explosives*. Retrieved July 29, 2021, from https://www.atf.gov/rules-and-regulations/bump-stocks

Bandini, S., Manzoni, S., & Vizzari, G. (2009). Agent Based Modeling and Simulation: An Informatics Perspective. *Journal of Artificial Societies and Social Simulation*, *12*(4), 4.

Bankes, S. C. (1992). *Exploratory Modeling and the Use of Simulation for Policy Analysis*. Rand. https://www.rand.org/content/dam/rand/pubs/notes/2009/N3093.pdf

Berner, C., Brockman, G., Chan, B., Cheung, V., Dębiak, P., Dennison, C., Farhi, D., Fischer, Q., Hashme, S., Hesse, C., Józefowicz, R., Gray, S., Olsson, C., Pachocki, J., Petrov, M., Pinto, H. P. de O., Raiman, J., Salimans, T., Schlatter, J., … Zhang, S. (2019). Dota 2 with Large Scale Deep Reinforcement Learning. *ArXiv:1912.06680 [Cs, Stat]*. http://arxiv.org/abs/1912.06680

Blair, J. P., & Martaindale, M. H. (2019). The Evolution of Active Shooter Response Training Protocols Since Columbine: Lessons From the Advanced Law Enforcement Rapid Response Training Center. *Journal of Contemporary Criminal Justice*, *35*(3), 342–356. https://doi.org/10.1177/1043986219840237

Blair, J. P., Nichols, T., Burns, D., & Curnutt, J. R. (2013). *Active Shooter Events and Response*. Routledge. https://doi.org/10.1201/b14996

Blair, J. Pete & Schweit, Katherine W. (2014). *A Study of Active Shooter Incidents in the United States Between 2000 and 2013*. Texas State University and Federal Bureau of Investigation, U.S. Department of Justice, Washington D.C. 2014. https://www.fbi.gov/file-repository/active-shooter-study-2000-2013-1.pdf/view

Bott, R. E., Frantz, B., & Dietz, E. J. (n.d.). Active Shooter Mitigation Techniques. *Journal of Emergency Management*, *In Press*, 33. https://doi.org/10.5055/jem.0639

Bradbury, S. (2021, March 27). Boulder shooting timeline: A chronology of the mass killing at King Soopers. *The Denver Post*. https://www.denverpost.com/2021/03/27/boulder-shooting-timeline-king-soopers-mass-killing/

BRIGGS CORE DYNAMICS. (n.d.). *LE Active Shooter*. BRIGGS CORE DYNAMICS. Retrieved June 25, 2021, from https://briggscoredynamics.com/le-active-shooter

Briggs, T. W., & Kennedy, W. G. (2016). Active shooter: An agent-based model of unarmed resistance. *2016 Winter Simulation Conference (WSC)*, 3521–3531. https://doi.org/10.1109/WSC.2016.7822381

Broward Sheriff's Office. (2018, April 25). *Nickolas Cruz Movements Through Marjory Stoneman Douglas*. YouTube. https://www.youtube.com/watch?v=bjqUejl4l58

Chen, C.-T., Chen, A.-P., & Huang, S.-H. (2018). Cloning Strategies from Trading Records using Agent-based Reinforcement Learning Algorithm. *2018 IEEE International Conference on Agents (ICA)*, 34–37. https://doi.org/10.1109/AGENTS.2018.8460078

City of Houston. (2012). *RUN.HIDE.FIGHT*. Run.Hide.Fight | Ready Houston. https://readyhouston.wpengine.com/suspicious-activity/run-hide-fight/

Cline, T., & Dietz, J. (2020). Agent Based Modeling for Low-cost Counter UAS Protocol in Prisons. *International Journal of Aviation, Aeronautics, and Aerospace*. https://doi.org/10.15394/ijaaa.2020.1462

Colloff, P. (2006, August). 96 Minutes. *Texas Monthly, August 2006*. https://www.texasmonthly.com/articles/96-minutes/

Corcoran, K., Baker, S., & Choi, D. (2019, January 29). The FBI has closed its investigation of the Las Vegas mass shooting that killed 58 people and injured hundreds more. Here's exactly how the nation's worst modern gun massacre unfolded [News]. *Business Insider*. https://www.businessinsider.com/timeline-shows-exactly-how-the-las-vegas-massacre-unfolded-2018-9

DeNolf, R. L., & Kahwaji, C. I. (2020). EMS Mass Casualty Management. In *StatPearls*. StatPearls Publishing. http://www.ncbi.nlm.nih.gov/books/NBK482373/

Editors, H. com. (n.d.). *Columbine Shooting*. HISTORY. Retrieved July 5, 2021, from https://www.history.com/topics/1990s/columbine-high-school-shootings

Editors, History. com. (2019, February 6). *Teen gunman kills 17, injures 17 at Parkland, Florida high school*. History. https://www.history.com/this-day-in-history/parkland-marjory-stoneman-douglas-school-shooting

El-Tantawy, S., & Abdulhai, B. (2010). An agent-based learning towards decentralized and coordinated traffic signal control. *13th International IEEE Conference on Intelligent Transportation Systems*, 665–670. https://doi.org/10.1109/ITSC.2010.5625066

Ervin Teng. (2019, November 11). Training your agents 7 times faster with ML-Agents. *Unity Blog*. https://blogs.unity3d.com/2019/11/11/training-your-agents-7-times-faster-with-ml-agents/

Farhan, M., Gohre, B., & Junprung, E. (2020). Reinforcement Learning in Anylogic Simulation Models: A Guiding Example Using Pathmind. *2020 Winter Simulation Conference (WSC)*, 3212–3223. https://doi.org/10.1109/WSC48552.2020.9383916

Federal Bureau of Investigation. (2021). *Active Shooter Incidents 20-Year Review, 2000-2019*. U. S. Department of Justice, Federal Bureau of Investigation. https://www.fbi.gov/file-repository/active-shooter-incidents-20-year-review-2000-2019-060121.pdf/view

Fleming, C. (2019, July 24). Remote Drone Dispatch: Law Enforcement's Future? *Police Chief Magazine*. https://www.policechiefmagazine.org/remote-drone-dispatch/

Frantz, B. M. (2021). *Active Shooter Mitigation for Open-Air Venues* [Purdue University]. https://doi.org/10.25394/PGS.14593671.v1

Géron, A. (2019). *Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow: Concepts, Tools, and Techniques to Build Intelligent Systems* (2nd ed.). O'Reilly Media. https://books.google.com/books?id=HHetDwAAQBAJ

Gross, D. (1999). *Fidelity Implementation Study Group* (Workshop No. 99S-SIW-167; Spring Simulation Interoperability Workshop). Simulation Interoperability Standards Organization. https://www.sisostds.org/DesktopModules/Bring2mind/DMX/API/Entries/Download?Command=Core_Download&EntryId=32793&PortalId=0&TabId=105

Hayes, R., & Hayes, R. (2014). Agent-Based Simulation of Mass Shootings: Determining How to Limit the Scale of a Tragedy. *The Journal of Artificial Societies and Social Simulation*, *17*(2), 1–12.

Jaclyn Schildkraut, & Elsass, H. J. (2021). Mass Shooting Factsheet. *Rockefeller Institute of Government*. https://rockinst.org/gun-violence/mass-shooting-factsheet/

Jalalimanesh, A., Shahabi Haghighi, H., Ahmadi, A., & Soltani, M. (2017). Simulation-based optimization of radiotherapy: Agent-based modeling and reinforcement learning. *Mathematics and Computers in Simulation*, *133*, 235–248. https://doi.org/10.1016/j.matcom.2016.05.008

Jennifer Campbell-Hicks. (2021, March 26). *"They put their lives at risk": Officers among first to respond to Boulder mass shooting saved others*. KUSA.Com. https://www.9news.com/article/news/local/boulder-shooting/boulder-shooting-king-soopers-police-update/73-5332fc5e-acf1-4648-bced-cef42490e5a9

Jiang, C. X., Jing, Z. X., Cui, X. R., Ji, T. Y., & Wu, Q. H. (2018). Multiple agents and reinforcement learning for modelling charging loads of electric taxis. *Applied Energy*, *222*, 158–168. https://doi.org/10.1016/j.apenergy.2018.03.164

Joel F. Shults. (2019, January 11). Parkland shooting report: Failures in communication and coordination. *Police1*. https://www.police1.com/mass-casualty/articles/parkland-shooting-report-failures-in-communication-and-coordination-CkeHCKIUbEx1VYTM/

Johnson, L., McBride, A., Fracassa, D., & Hernández, L. (2019, July 29). How the Gilroy Garlic Festival shooting unfolded. *San Francisco Chronicle*. https://www.sfchronicle.com/crime/article/Gilroy-Garlic-Festival-shooting-Violence-14193120.php

Jonson, C. L., Moon, M. M., & Hendry, J. A. (2020). One Size Does Not Fit All: Traditional Lockdown Versus Multioption Responses to School Shootings. *Journal of School Violence, 19*(2), 154–166. https://doi.org/10.1080/15388220.2018.1553719

Joseph Lombardo. (2018). *LVMPD Criminal Investigative Report of the 1 October Mass Casualty Shooting* (Criminal Investigative Report No. 171001–3519). https://www.lvmpd.com/en-us/Documents/1-October-FIT-Criminal-Investigative-Report-FINAL_080318.pdf

Kiran, P., & Chandrakala, V. (2020). New interactive agent based reinforcement learning approach towards smart generator bidding in electricity market with micro grid integration. *Applied Soft Computing*, *97*, 106762. https://doi.org/10.1016/j.asoc.2020.106762

Kirby, A., Anklam, C. E., & Dietz, J. E. (2016). Active shooter mitigation for gun-free zones. *2016 IEEE Symposium on Technologies for Homeland Security (HST)*, 1–6. https://doi.org/10.1109/THS.2016.7568957

Kirby, A., Dietz, J. E., & Wojtalewicz, C. (2012). Modeling of a Regional Hub Reception Center to improve the speed of an urban area evacuation. *2012 IEEE Conference on Technologies for Homeland Security (HST)*, 476–482. https://doi.org/10.1109/THS.2012.6459895

Kirby, A. M. (2016). COMPARING POLICY DECISIONS FOR ACTIVE SHOOTERS USING SIMULATION MODELING. *ProQuest*, 24.

Knox, M. (2018). *Crime Scene Behaviors of Rampage School Shooters: Developing Strategies for Planning, Response, and Investigation of Multiple-Victim Shooting Incidents on School Campuses* [Dissertation, Nova Southeastern University]. https://nsuworks.nova.edu/cgi/viewcontent.cgi?article=1006&context=cahss_jhs_etd

Kristopher D. Davis, Braiden M. Frantz, Jae Y. Lee, Zachery N. Panzer, & Qingyan Zhang. (n.d.). Validating and Optimizing an Agent-Based Simulation Model to Improve Law Enforcement Active Shooter Training. *Unpublished*, 22.

Krouse, W. J., & Richardson, D. J. (2015). Mass Murder with Firearms: Incidents and Victims, 1999-2013. *Congressional Research Service*, 51.

Larsen, E. (2021, March 23). *Arrest documents for Boulder shooting suspect detail how deadly rampage unfolded*. The Coloradoan. https://www.coloradoan.com/story/news/2021/03/23/boulder-shooting-timeline-shooter-ahmad-alissa-king-soopers-police/6968055002/

Lee, J. Y. (2019). *AGENT-BASED MODELING TO ASSESS THE EFFECTIVENESS OF RUN HIDE FIGHT* [Thesis, Purdue University Graduate School]. https://doi.org/10.25394/PGS.8020763.v1

Macal, C. M., & North, M. J. (2005). Tutorial on agent-based modeling and simulation. *Proceedings of the Winter Simulation Conference, 2005.*, 2–15. https://doi.org/10.1109/WSC.2005.1574234

Mallonee, L. (2017, September 5). Inside Texas' Active Shooter Training Simulations. *Wired*. https://www.wired.com/story/inside-texas-active-shooter-training-simulations/

Marcou, D. (2019, October 1). After-action review: Route 91 Harvest Music Festival mass shooting. *Police1*. https://www.police1.com/active-shooter/articles/after-action-review-route-91-harvest-music-festival-mass-shooting-Ujz7O1z77ibVGb4F/

Mattes, P. W. (2019, October 1). What is a Modern Integrated Air Defense System. *Air Force Magazine*. https://www.airforcemag.com/article/what-is-a-modern-integrated-air-defense-system/

McBride, J. (2021, March 22). Live Graphic Video: King Soopers Active Shooter in Boulder, Colorado. *Heavy.Com*. https://heavy.com/news/king-soopers-live-video-active-shooter-boulder/

MSD Public Safety Commission. (2019a). *MARJORY STONEMAN DOUGLAS HIGH SCHOOL PUBLIC SAFETY COMMISSION INITIAL REPORT*. http://www.fdle.state.fl.us/MSDHS/CommissionReport.pdf

MSD Public Safety Commission. (2019b). *MARJORY STONEMAN DOUGLAS HIGH SCHOOL PUBLIC SAFETY COMMISSION SECOND REPORT*. http://www.fdle.state.fl.us/MSDHS/MSD-Report-2-Public-Version.pdf

National Rifle Association of America. (n.d.). *NRA School Shield*. NRA School Shield. Retrieved October 26, 2020, from https://www.nraschoolshield.org/

NBC New York. (2019, July 28). *Mass Shooting Reported at Gilroy Garlic Festival in California | NBC New York*. YouTube. https://www.youtube.com/watch?v=XNnmLcPyYWk

Okoro, D. (2021, July 7). *Slain Dallas officers remembered on 5-year anniversary of sniper attack*. https://spectrumlocalnews.com/tx/south-texas-el-paso/news/2021/07/07/slain-dallas-officers-remembered-on-5-year-anniversary-of-sniper-attack

O'Matz, M., Fleshler, D., & Alanez, T. (2018, November 14). Deputies who responded to Parkland tragedy hadn't had active shooter training in years, investigators found. *Sun-Sentinel.Com*. https://www.sun-sentinel.com/local/broward/parkland/florida-school-shooting/fl-ne-msd-police-training-response-20181114-story.html

OpenAI. (n.d.-a). *About OpenAI*. Retrieved March 17, 2020, from https://openai.com/about/

OpenAI. (n.d.-b). *Gym*. OpenAI. Retrieved March 18, 2020, from https://gym.openai.com/

Paris Buttfield-Addison. (2019, November 5). *Let's build a video game, train a bot to play it, and deploy it on a smartphone in 30 minutes*. North Bay Python 2019. https://www.youtube.com/watch?v=8fLyMPPixdU

Pathmind Inc. (2021). *Simulation Optimization | Add AI to Simulation Models*. Pathmind Inc. https://pathmind.com/

Phillips, S. W. (2020). Police Response to Active Shooter Events: How Officers See Their Role. *Police Quarterly*, *23*(2), 262–279. https://doi.org/10.1177/1098611119896654

Piao, H., Sun, Z., Meng, G., Chen, H., Qu, B., Lang, K., Sun, Y., Yang, S., & Peng, X. (2020). Beyond-Visual-Range Air Combat Tactics Auto-Generation by Reinforcement Learning.

2020 International Joint Conference on Neural Networks (IJCNN), 1–8. https://doi.org/10.1109/IJCNN48605.2020.9207088

Police1. (2007, August 31). The Stopwatch of Death. *Police1*. https://www.police1.com/police-products/police-technology/emergency-response/articles/the-stopwatch-of-death-SYnKshr9PnUurx3Y/

Rasmussen, C. T. E., Baer, D. D. G., Doll, R. B. A., & Jr, M. J. C. (2015). *Combat Casualty Care Research drives innovation to improve survivability and reimagine future combat care*. 6.

Ray, M. (2020, December 7). *Sandy Hook Elementary School shooting | Facts & Timeline*. Encyclopedia Britannica. https://www.britannica.com/event/Newtown-shootings-of-2012

Rosen, J. (2020). *Report on the Non-Fatal Shooting of Santino Legan on July 28, 2019* (p. 26) [Officer Involved Shooting]. Santa Clara County. https://www.sccgov.org/sites/da/newsroom/newsreleases/Documents/2020NRDocs/Santino%20Legan.pdf

Rubino, J. (2021, March 22). Boulder shooting: Gunman kills 10, including police officer, at King Soopers. *The Denver Post*. https://www.denverpost.com/2021/03/22/police-active-shooter-shooting-king-soopers-boulder/

Saracco, R. (2019, July 28). Drones growing market. *IEEE Future Directions*. https://cmte.ieee.org/futuredirections/2019/07/28/drones-growing-market/

Sauer, P. (2015, October 14). The Story of the First Mass Murder in U.S. History. *Smithsonian Magazine*. https://www.smithsonianmag.com/history/story-first-mass-murder-us-history-180956927/

Schildkraut, J., Formica, M. K., & Malatras, J. (2018, May 22). Can Mass Shootings be Stopped? *Rockefeller Institute of Government*. https://rockinst.org/issue-area/can-mass-shootings-be-stopped/

Silva, J. R., & Capellan, J. A. (2019). The Media's Coverage of Mass Public Shootings in America: Fifty Years of Newsworthiness. *International Journal of Comparative and Applied Criminal Justice*, *43*(1), 77–97. https://doi.org/10.1080/01924036.2018.1437458

Silver, D., & Hassabis, D. (2017, October 18). *AlphaGo Zero: Starting from scratch*. https://deepmind.com/blog/article/alphago-zero-starting-scratch

Smart, R., & Schell, T. L. (2021, April 15). *Mass Shootings in the United States*. https://www.rand.org/research/gun-policy/analysis/essays/mass-shootings.html

Somers, J. (2018, December 28). How the Artificial-Intelligence Program AlphaZero Mastered Its Games. *The New Yorker*. https://www.newyorker.com/science/elements/how-the-artificial-intelligence-program-alphazero-mastered-its-games

Stephen J. Sedensky III. (2013). *Report of the State's Attorney for the Judicial District of Danbury on the Shootings at Sandy Hook Elementary School and 36 Yogananda Street, Newtown, Connecticut on December 14, 2012* (p. 44). OFFICE OF THE STATE'S ATTORNEY JUDICIAL DISTRICT OF DANBURY. https://portal.ct.gov/-/media/DCJ/SandyHookFinalReportpdf.pdf

Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction* (Second edition). The MIT Press.

Texas State University. (n.d.). *AvoidDenyDefend*. AvoidDenyDefend.Org. Retrieved June 25, 2021, from https://www.avoiddenydefend.org

Texas State University. (2021). *Advanced Law Enforcement Rapid Response Training*. Advanced Law Enforcement Rapid Response Training. https://alerrt.org/

The AlphaStar team. (2019, January 24). *AlphaStar: Mastering the Real-Time Strategy Game StarCraft II*. https://deepmind.com/blog/article/alphastar-mastering-real-time-strategy-game-starcraft-ii

Tzimas, A., Passalis, N., & Tefas, A. (2020). Leveraging Deep Reinforcement Learning For Active Shooting Under Open-World Setting. *2020 IEEE International Conference on Multimedia and Expo (ICME)*, 1–6. https://doi.org/10.1109/ICME46284.2020.9102966

United States Department of Justice Office for Victims of Crimes. (2017, January 1). 2017 National Crime Victims' Rights Week Resource Guide: Crime and Victimization Fact Sheets: Mass Casualty Shootings. *Homeland Security Digital Library*. National Center for Victims of Crime (U.S.); United States. Office of Justice Programs. Office for Victims of Crime. https://www.hsdl.org/?view&did=819184

Unity. (n.d.-a). *Unity ml-agents*. Unity. Retrieved October 22, 2020, from https://unity.com/products/machine-learning-agents

Unity. (n.d.-b). *Unity ML-Agents Gym Wrapper*. Unity. https://github.com/Unity-Technologies/ml-agents/tree/master/gym-unity

Virginia Tech Review Panel. (2007). *Mass Shootings at Virginia Tech* (pp. 1–146) [Review Panel]. Virginia Tech University.

What is Reinforcement Learning? – Overview of How it Works | Synopsys. (2021, April 27). https://www.synopsys.com/ai/what-is-reinforcement-learning.html

Wilensky, U. (1999). *NetLogo*. Center for Connected Learning and Computer-Based Modeling, Northwestern University. Evanston, IL. https://ccl.northwestern.edu/netlogo/

Xia, S.-M., Guo, S.-Z., Bai, W., Qiu, J.-Y., Wei, H., & Pan, Z.-S. (2019). A New Smart Router-Throttling Method to Mitigate DDoS Attacks. *IEEE Access*, *7*, 107952–107963. https://doi.org/10.1109/ACCESS.2019.2930803

ZFG Videography. (2021, March 22). *First on Scene at Mass Shooting, King Soopers Boulder CO* [Video]. YouTube. https://www.youtube.com/watch?v=ikujTMetrUk