Evacuation in the Presence of Bad Actors

Clark Petri[1] and Hamdi Kavak[1]

George Mason University, Fairfax VA 22030, USA

Abstract. In the wake of the January 6th insurrection and 2020 protests, modeling the tipping point of crowds facing adversarial entities has become a more timely and pertinent research field. Limited research has been performed on the effects of discreet malevolent agents compared to overt bad actors such as mass shooters or suicide bombers. Furthermore, the current literature on evacuation models does not consider the presence of bad actors. In this paper, we build an abstract evacuation model where agents operating in a cellular space proceed toward an 'exit' location across two scenarios: one without a malevolent actor(s) and the other with one. In the latter scenario, the bad actor(s) will reach the exit but not leave, impeding efficient evacuation. Analysis shows that crowd agitation levels, the presence of individuals willing to engage physically, and the loss of an exit drive total evacuation time.

Keywords: agent-based modeling · bad actors · crowd dynamics · evacuation · MESA.

1 Introduction

This paper explores the impact of malevolent agents that seek to disrupt an evacuation and the resulting crowd dynamics initiating conflict between the evacuating and malevolent agents. Compared to the modeling efforts of other bad actors that are explicitly lethal, i.e., active shooters or suicide bombers, the malevolence of agents impeding efficient evacuation is limited. Still, the combination of evacuation and subversion within a crowd dynamic framework, to include modeling the tipping point of violence, is novel Agent-Based Modeling (ABM) research.

Our model sits at the intersection of three research domains: evacuation modeling, cooperative agents/crowds, and bad actor modeling. This is illustrated in figure 1 with the pertinent literature annotated by domain. It is the unique combination of topics that makes our work novel within the field. Figure 1 serves as the structure for the following literature review.

1.1 Evacuation Modeling

There are many examples of evacuation scenario research and modeling, frequently for the purposes of fire evacuation. Owen et al. [14] developed a building fire evacuation model, representing a hypothetical supermarket and taking

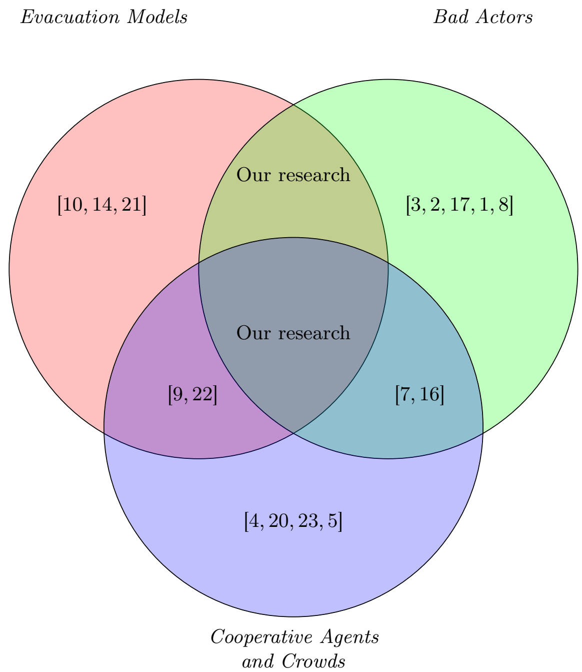

Fig. 1. Synthesis of evacuation, cooperation / crowds, and bad actor related literature.

agent characteristics such as movement speed, size, and reaction times into account. Their results show a variety of evacuation performance measures based on changes in pedestrian travel speed and response times.

Hou et al. produced a social force evacuation model evaluating the efficacy of evacuation leaders and their presence near exits, concluding that emergency plans should identify leaders for the number of exits necessary in a given scenario [10]. The use of cellular automata (CA) space, similar to the model presented in this paper, has been employed. For example, the work of Yuan and Tan [21] used CA to model fire evacuation and took visual spatial distance, occupant density, and human behavior effects into account.

1.2 Bad Actors

There is abundant research on the impact of overtly malevolent agents, ranging from their effect on other non-threat agents to the efficient functioning of systems. Briggs and Kennedy [7] produced an Agent-Based Model (ABM) for active shooter scenarios, specifically exploring the "run, hide, fight" resistance paradigm. The cooperative resistance elements of their research place their work in the union of Bad Actors and Cooperative Agents. Further analysis of overtly malevolent actors in active shooter scenarios has been performed by Arteaga and Park [3], as well as Anklam et al. [2]. The master's thesis by Stewart [17] specifically focused on classroom active shooter scenarios, with three primary parameters: law enforcement response time, civilian response strategy, and cognitive delay.

Another well-researched category of bad actors includes suicide bombers. Alghamdi et al. [1] put forward a general blast wave model accounting for crowd density, partial and full cover, and type/quantity of explosive. Other research includes suicide bomber target decision making by Bulleit and Drewek [8], their efforts having similarity to the classic linear programming attacker-defender or defender-attacker-defender ontologies.

1.3 Cooperative Agents

Copious research exists regarding cooperative agents. Axelrod and Hamilton [4] famously demonstrated the effectiveness of cooperation in the prisoner's dilemma, notably when tit-for-tat agents meet each other. Other examples include the benefits of cooperation via contribution with empirical inputs derived from observations of over 150 human players engaged in public good games by Wunder et al. [20]. Other domain-specific research showing the benefits of the cooperative game theory include Wang et. al. [18], as well as Zolezzi and Rudnick [23].

Reinforcement learning produces cooperation among agents as well, notably (and adorably) with Baker et al. [5] employing OpenAI's tools to multi-agent hide and seek. Cooperation emerged to produce increasingly complicated strategies for both the hiders as well as the seekers. The evolutionary back and forth mirrors the defender-attacker-defender paradigm, looped indefinitely, as cooperation manifested increasingly novel behavior of the agents.

Research at the intersection of evacuation and cooperative agents has been performed by Cheng et al. [9], finding that cooperation emerges in emergency scenarios when agents have high levels of escape aspiration. Other research demonstrated crowd imitation increases cooperation in evacuation scenarios, but higher degrees of emergency lower cooperation [22].

The distinction between cooperation/non-cooperation and bad actors is important. Self-serving non-cooperation in evacuation models can limit evacuation efficiency, as shown by Zheng and Cheng [22]. However, lessened efficiency via non-cooperative self-interest is distinctly different from active impedance of system function with bad actors.

Lastly but importantly, Reicher [16] put forward an analysis of the student protest cooperation and crowd dynamics during the "Battle of Westminster" when clashes between students and police seemingly erupted spontaneously. Across many interviews, Reicher found that agitation, the perception that a right is being deprived, and the presence of individuals willing to engage in the physical violence are the requirements necessary to spur a crowd to collective action. Reicher's research underpins the agent agitation and engagement logic in our model.

2 Model

2.1 Environment and Model Parameters

The environment of the model is designed based on a spatial grid with a 20x20 non-toroidal cells. We assume that only one agent may occupy a cell, and all agents are randomly placed for a given model iteration. For the purposes of evacuation and adding spatial components representing a room, non-interacting "wall" entities are added, structuring two doors leaving the space as well as a central pillar. We envision this as a conference room. Figure 2 shows an instance of the model space, with evacuation agents in blue, evacuation agent heroes in green, malevolent agents in red, wall entities in black, and the exits outlined in green.

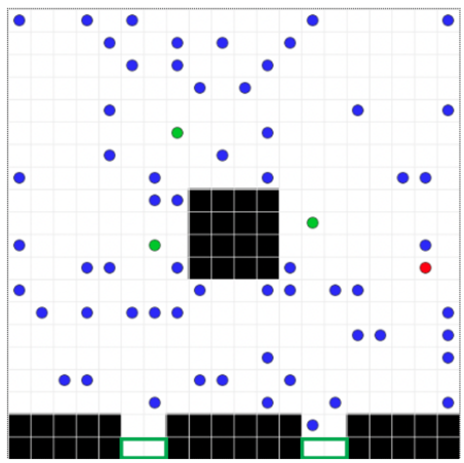

Fig. 2. Model space example with agent density = 0.2, 1 malevolent agent, and 3 hero agents.

Model parameters are *agent density* as a fraction of the total model space (variable range between 0.1 to 1.0), the *number of malevolent agents* ranging from 0 to 3, *the probability of an evacuation agent being a hero* ranging 0 to 0.5, *the tipping point threshold* for crowd violence on a 0 to 1 scale, and *the agitation increase* of evacuation agents near the exit ranging from 0 to 1.0 each model step. Default initialization is set to density=0.75, malevolent=1, hero probability at 2.5%, and a 0.5 tipping point. However, these settings can be changed within the browser-based graphical user interface (GUI) or in code for batch runs.

2.2 Agent Behavior

At each model step, each agent activates in random order with the order reshuffled each step. During activation, the agent calculates the shortest Euclidean distance to the exits for each available empty space around the agent as a Moore neighborhood [13]. Figure 2 shows the currently active agent identified in orange outline with available movement options a, b, c, d, e, f. The distance between each movement option and the nearest exits, denoted as x and y, are calculated. In the Figure 2 example, the c, x and d, y pairs have equal shortest distances. The agent will select the first sequential best move, in this case, moving to position c.

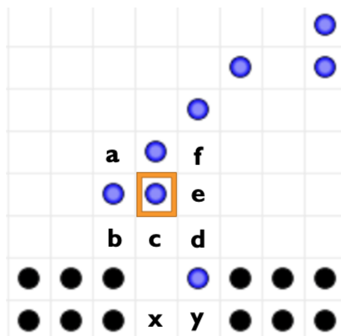

Fig. 3. Example agent movement options.

Provided at least one available move, agents will move each step, even if a move brings the agent further from an exit. Referring to the example in Figure 2, if the only available moves were locations a and f the agent would move to f. While suboptimal for evacuation, we contend that this behavior resembles crowd shuffling and the chaos of emergency scenarios.

Upon reaching an exit location, evacuation agents are removed from the model and schedule. The simulator logs these agents as part of an escaped count, as well as the total remaining evacuation agents and the current step. Malevolent agents exhibit identical behavior as evacuation agents, with the exception that upon reaching an exit, they are removed from the schedule but not the model, blocking the exit cell. The model run concludes when all evacuation agents have escaped.

The conflict between the evacuation and malevolent agents is represented by the removal of the malevolent agent and restoration of the exit, initiated when the following criteria are met: One, evacuation agents possess an aggregate agitation score above the tipping point threshold. Agitation increases each model step when the agent is within a 3 x 6 zone near an exit.

Second, a malevolent agent is required to be blocking an exit space. Third and finally, the evacuation zone requires the presence of a hero evacuation agent, one that is willing to "physically confront" the malevolent agent. With these criteria met the model's engagement parameter is set to "True" and the malevolent agent(s) impeding the exit(s) are removed.

2.3 Implementation

The model was built in Python 3.10.2 and predominantly structured in Mesa 0.9.0. Mesa "allows users to quickly create agent-based models using built-in core components (such as spatial grids and agent schedulers) or customized implementations; visualize them using a browser-based interface; and analyze their results using Python's data analysis tools" [11]. As a somewhat nascent Agent-Based Modeling (ABM) environment, Mesa does not natively include the ability to create irregular shaped grids, strictly allowing for m x n structures. The use of non-interacting "wall agents" is our work-around to this limitation which, to our knowledge, does not exist within other Mesa models.

3 Analysis of Results

3.1 Parameter Settings and Coefficient of Variance Determination

Mesa includes a BatchRunner module enabling multicore batch runs, parameter sweeps, and data output/capture for a given model [11]. Preliminary analysis was performed with data captured from the following parametric combinations, of which there are 2592 total. Replications ranging from 5 to 50 were performed for each parameter combination set. The parameter ranges were selected based on saving computational resources via limiting dimensional growth, informed by initial overview analysis and testing. For example, limited statistical change exists with *Hero Probability* at 0.50 and above due to continuous hero presence.

Table 1. Batch run settings.

Parameter	Setting(s)
Density	.25/.50/.75
Hero Probability	0/.10/.20/.30/40/.50
Malevolent Agents	0/1/2/3
Tipping Point	0/.20/.40/.60/.80/1.0
Agitation Gain	0/.20/.40/.60/.80/1.0

Due to the 2592 parametric combinations and the functional similarities between different sets of combinations, the coefficient of variance testing required few replications. We use the equation (1) proposed by Lee et al. [12]. Here, σ represents the standard deviation, and μ represents the mean. We found that the increase from 5 to 10 replications produced a coefficient of variance change under epsilon 0.01; thus, analysis in the following sections was performed on data output from the 10 replications batch run.

$$c_v = \frac{\sigma}{\mu} \tag{1}$$

3.2 Time to Evacuate

Total steps for all agents to evacuate is our primary metric of interest and is most impacted by agent density. The three density settings, 0.25, 0.50, and 0.75, are clearly visible via the clustering in Figure 4.

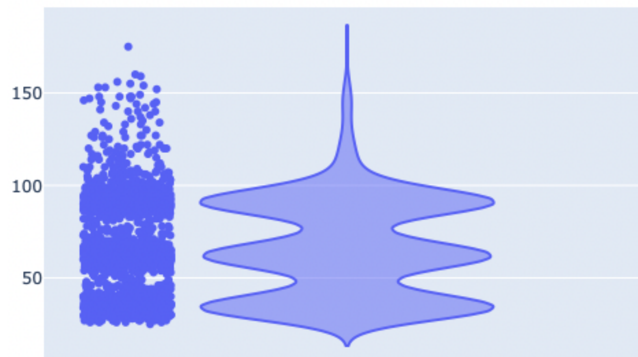

Fig. 4. Violin plot of total evacuation time in steps.

Scenarios with 2 or 3 malevolent agents, i.e., situations with the potential to fully lose access to an exit, are those showing the maximum time to complete evacuation and the largest standard deviation in total model steps. However, permanent loss of an exit was only possible in parametric combinations where the model's engagement parameter was precluded from switching to *True*, requiring one of the following:

– Hero probability equal to 0.0
– Tipping point of 1.0
– Agitation gain of 0.0

Without one of the above parameter settings, the potential for engagement exists, and the loss of an exit is frequently temporary. For simulations where engagement was possible, i.e., none of the above parameter settings and at least one malevolent agent, engagement occurred 74.6% of the time across the remaining parameter space.

Model behavior in terms of time to evacuate and the conditions affecting evacuation time the most are as expected, i.e., more total agents, more malevolent agents, and parameter settings lowering the likelihood of engagement correlate with longer evacuation time on average. Due to this aggregate behavior as intended and anticipated, we conclude the model is verified.

3.3 Engagement Analysis

The factors and conditions leading to changes in crowd behavior are of broader and continuing research interest. Using the data obtained from the batch runs, we conducted Logistic Regression, and machine-learning [15] based feature importance measures. The five parameters affecting the potential for engagement (density, hero probability, number of malevolent agents, tipping point, and agitation increase) were set as independent variables with the True/False boolean for engagement as the dependent variable.

The logistic regression model was trained after an extensive hyperparameter optimization leading to a maximum accuracy score of 0.77. This result tells us that a linear model can explain relatively high variation; however, there are relationships that require a non-linear model. To go deeper in our evaluation, we looked at the Logistic Regression's coefficients as shown in Table 2.

Table 2. Logistic Regression Results

Parameter	Coefficient
Density	-0.16
Hero Probability	0.77
Malevolent Agents	1.16
Tipping Point	-0.95
Agitation Gain	0.78

Here, except for the *density* parameter's coefficient, the remaining four parameters' coefficients were statistically significant. However, their contribution was not all identical. The coefficient with the highest contribution was the number of *Malevolent Agents*, which impacts engagement positively. *Hero Probability* and *Agitation Gain* were the other two positive parameters that contribute to the odds of having an engagement. These positive contributions make sense and are expected. The *Tipping Point* parameter's coefficient, on the other hand, negatively contributed to the odds of having an engagement. This also makes sense because a higher tipping point will lead to less engagement.

We have also created a Random Forest [6] classifier to verify the coefficients found in the Logistic Regression model. With the hyperparameter optimization, the model performed approximately 99.5% accuracy and provided a Gini-based [6] feature importance measure. The results we found confirmed our previous co-efficient finding.

4 Conclusions and Future Research

We successfully implemented an evacuation scenario within the Mesa environment and crafted a solution for irregularly shaped cellular automata that, to our knowledge, do not exist within current Mesa models. Model runs and analysis demonstrate the importance of maintaining maximum exit availability in the event of an evacuation, with the complete loss of an exit being the largest determinant of extended evacuation time (steps), aside from initial density. While expanding agent behavior and shaping against actual locales and known evacuation data is necessary for model verification, we conclude the model is validated due to its behavior.

Our exploration of crowd behavior in the presence of malevolent actors is potentially the first in a new research domain. Combining the techniques of other evacuation models with our preliminary work will produce further insights into system performance in bad actors' presence. Moreover, continued work within a Python and Mesa environment will empower the use of more robust analytical tools and feedback within a single development framework, as opposed to the post-simulation limitations presented by established tools such as NetLogo [19].

We are specifically focused on two focal points for further development: improved crowd behavior and real location modeling. Expanded crowd/agent behavior includes implementing a social force model of movement or distinct agent movement qualities. Examples include the line of sight and variable movement speeds. Implementing cooperation among malevolent agents or the ability to fight back when engaged are other areas for refinement.

Adapting the model to simulate real evacuation scenarios is our other focal point for continued research. Comparisons of simulated results to reality are steps toward model verification.

Notes: Additional information, code, and datasets for this paper are freely available at https://github.com/clarkpetri/evacuation-model.

Acknowledgements: Clark Petri would like to thank his wife Sherri for her love, strength, and support, as well as Dr. Kavak for putting up with his neverending questions.

References

1. Alghamdi, F.A., Kirk, D., et al.: Blastsim—multi agent simulation of suicide bombing. In: 2009 IEEE Symposium on Computational Intelligence for Security and Defense Applications. pp. 1–8. IEEE (2009)
2. Anklam, C., Kirby, A., Sharevski, F., Dietz, J.E.: Mitigating active shooter impact; analysis for policy options based on agent/computer based modeling. Journal of Emergency Management (2014)
3. Arteaga, C., Park, J.: Building design and its effect on evacuation efficiency and casualty levels during an indoor active shooter incident. Safety science **127**, 104692 (2020)
4. Axelrod, R., Hamilton, W.D.: The evolution of cooperation. science **211**(4489), 1390–1396 (1981)
5. Baker, B., Kanitscheider, I., Markov, T., Wu, Y., Powell, G., McGrew, B., Mordatch, I.: Emergent tool use from multi-agent autocurricula. arXiv preprint arXiv:1909.07528 (2019)

6. Breiman, L.: Random forests. Machine learning **45**(1), 5–32 (2001)
7. Briggs, T.W., Kennedy, W.G.: Active shooter: An agent-based model of unarmed resistance. In: 2016 Winter Simulation Conference (WSC). pp. 3521–3531. IEEE (2016)
8. Bulleit, W.M., Drewek, M.W.: An agent-based model of terrorist activity. Proceedings of the North American Association for Computational Social and Organizational Science (NAACSOS 2005) (2005)
9. Cheng, Y., Zheng, X.: Emergence of cooperation during an emergency evacuation. Applied Mathematics and Computation **320**, 485–494 (2018)
10. Hou, L., Liu, J.G., Pan, X., Wang, B.H.: A social force evacuation model with the leadership effect. Physica A: Statistical Mechanics and its Applications **400**, 93–99 (2014)
11. Kazil, J., Masad, D., Crooks, A.: Utilizing python for agent-based modeling: The mesa framework. In: Thomson, R., Bisgin, H., Dancy, C., Hyder, A., Hussain, M. (eds.) Social, Cultural, and Behavioral Modeling. pp. 308–317. Springer International Publishing, Cham (2020)
12. Lee, J.S., Filatova, T., Ligmann-Zielinska, A., Hassani-Mahmooei, B., Stonedahl, F., Lorscheid, I., Voinov, A., Polhill, J.G., Sun, Z., Parker, D.C.: The complexities of agent-based modeling output analysis. Journal of Artificial Societies and Social Simulation **18**(4) (2015)
13. Moore, E.F., et al.: Gedanken-experiments on sequential machines. Automata studies **34**, 129–153 (1956)
14. Owen, M., Galea, E.R., Lawrence, P.J.: The exodus evacuation model applied to building evacuation scenarios. Journal of Fire Protection Engineering **8**(2), 65–84 (1996)
15. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E.: Scikit-learn: Machine learning in Python. Journal of Machine Learning Research **12**, 2825–2830 (2011)
16. Reicher, S.D.: 'the battle of westminster': Developing the social identity model of crowd behaviour in order to explain the initiation and development of collective conflict. European journal of social psychology **26**(1), 115–134 (1996)
17. Stewart, A.: Active shooter simulations: An agent-based model of civilian response strategy. Ph.D. thesis, Iowa State University (2017)
18. Wang, Z., Wu, G., Boriboonsomsin, K., Barth, M.J., Han, K., Kim, B., Tiwari, P.: Cooperative ramp merging system: Agent-based modeling and simulation using game engine. SAE International Journal of Connected and Automated Vehicles **2**(2) (2019)
19. Wilensky, U.: Netlogo. http://ccl.northwestern.edu/netlogo/, Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL (1999), http://ccl.northwestern.edu/netlogo/
20. Wunder, M., Suri, S., Watts, D.J.: Empirical agent based models of cooperation in public goods games. In: Proceedings of the fourteenth ACM conference on electronic commerce. pp. 891–908 (2013)
21. Yuan, W., Tan, K.H.: An evacuation model using cellular automata. Physica A: Statistical Mechanics and its Applications **384**(2), 549–566 (2007)
22. Zheng, X., Cheng, Y.: Modeling cooperative and competitive behaviors in emergency evacuation: A game-theoretical approach. Computers & Mathematics with Applications **62**(12), 4627–4634 (2011)
23. Zolezzi, J.M., Rudnick, H.: Transmission cost allocation by cooperative games and coalition formation. IEEE Transactions on power systems **17**(4), 1008–1015 (2002)