

Studying the Effects of Robot Distraction on School Shooter Behavior Using Virtual Reality

Christopher A. McClurg, Alan R. Wagner

The Pennsylvania State University

Abstract

We examine a particular adversarial human–robot interaction for which mobile robots influence the behavior of people role-playing as a school shooter. Using a high-fidelity virtual reality environment, we conducted a controlled study with 150 participants. Two autonomous robots predicted participant movement and positioned themselves to interfere and distract. The robots’ approach strategy was manipulated—either moving directly into the participant’s path (aggressive) or maintaining distance (passive)—along with the level of distraction, ranging from no additional cues (low), to siren and lights (medium), to siren, lights, and smoke to impair visibility (high). An aggressive, high-distraction robot configuration reduced the number of victims by 46.6% relative to a no-robot control. These results show that robot-based distraction can meaningfully alter human behavior and outcomes in adversarial settings, while also raising important ethical questions about the use of such systems in school environments.

1 Introduction

Regrettably, school shootings in the United States have increased dramatically. Among incidents involving at least one victim, there were more school shootings in the last decade (2016–2025) than in the previous five decades combined (1966–2015) [Riedman, 2025]. The harm caused by school shootings goes beyond the immediate victims. Students who witness a shooting suffer high rates of Post-Traumatic Stress Disorder (PTSD) and psychiatric disorders [Elklit and Kurdahl, 2013; Suomalainen *et al.*, 2011].

Various approaches have been implemented to prevent school shootings, including hardening school buildings, adding security measures such as metal detectors, zero-tolerance bullying programs, and profiling potential active shooters [Mongan *et al.*, 2009]. The effectiveness of these efforts is unclear [Jonson, 2017; Crawford and Burns, 2015; Borum *et al.*, 2010]. More than half of schools have hired armed school resource officers [Jonson, 2017; Addington, 2009; U.S. Department of Health and Human Services, 2015], even though armed officers are costly [Addington,



Figure 1: This study examines human behavior during a role-playing task in which participants, acting as a shooter, interact with mobile robots using loud noises, flashing lights, and/or smoke to disrupt attention and decision-making.

2009], may signal to students that schools are unsafe [Bachman *et al.*, 2011], and may not be an effective deterrent [James and McCallion, 2013].

This paper examines¹ whether a mobile robot can delay or divert the attention of an active shooter, as illustrated in Fig. 1. Rather than deploying a robot specifically for shooter intervention, we envision that robots already present in the school for other purposes, such as cleaning, could be activated during an emergency by school alarms or local emergency personnel. Once activated, the robot would locate the shooter using existing surveillance cameras, estimate the shooter’s future location, and move toward the shooter to initiate a distraction. We focus on the latter stages of this scenario, namely the effectiveness of the robot’s strategy for approaching the shooter and its method of distraction. Our overarching objective is to understand how to design a robot that could one day delay an active shooter in order to save lives.

We assume that the robot is notified of the shooter’s presence and relies on existing surveillance infrastructure to estimate the shooter’s location. Prior work has demonstrated the feasibility of tracking people during simulated emer-

¹This study was approved by an institutional review board (IRB) prior to data collection.

agencies [Nayyar and Wagner, 2022; Nayyar and Wagner, 2023]. We further assume that students and faculty shelter in place, as recommended and observed in prior school shootings [Ausdemore, 2015], and that the school environment is mapped such that a robot can localize and navigate using standard sensing and robotics frameworks [Quigley *et al.*, 2009]. Under these assumptions, a robot can navigate to the general vicinity of the shooter.

As a step toward understanding how robots might influence human behavior in violent emergencies, we conducted a human-subject study in virtual reality in which university students and faculty role-played as a school shooter during a simulated incident. The study compared two robot approach strategies and three methods of distraction to evaluate their effects on shooter behavior and impact on the number of victims. Our results provide empirical evidence that robotic distraction can delay shooter progress and reduce the number of victims in this setting. To the best of our knowledge, this is the first study to examine adversarial human-robot interaction in the context of a school shooting using controlled human-subject experiments.

2 Related Work

For this research, the human shooter and the robot are adversaries. We therefore review prior work on adversarial robotics, focusing on systems developed for human adversaries and the extent to which human behavior is evaluated.

Non-lethal adversarial robots have been developed for law enforcement and patrol applications [Nguyen and Bott, 2001; Lundberg and Christensen, 2007; Knutsson, 2025; Cramer and Hauser, 2021]. A survey of law enforcement agencies indicated that, if available, robots would be used frequently to inspect hazardous areas or as a tool for hostage negotiation [Nguyen and Bott, 2001]. In another study, a special weapons and tactics unit (SWAT) was given a mobile robot with camera and siren for five months and concluded that the robot was suitable for hostage negotiation [Lundberg and Christensen, 2007]. Law enforcement in China has introduced a 275-pound spherical robot capable of speeds up to 22 miles per hour to track and pursue suspects [Knutsson, 2025]. In the United States, law enforcement has tested a four-legged robot dog for similar purposes, where it has shown promise for use in hostage situations [Cramer and Hauser, 2021]. Much of this work emphasizes feasibility and operational effectiveness from the perspective of law enforcement, with limited empirical analysis of human adversary behavior.

Adversarial robot patrolling. Whereas regular patrolling focuses on area coverage and threat detection, adversarial patrolling assumes the presence of an intelligent adversary [Huang *et al.*, 2019]. Agmon *et al.* introduced a multi-robot patrolling strategy based on randomized Markovian policies to improve robustness against learned adversaries [Agmon *et al.*, 2011]. Subsequent work has examined coordinated attacks [Sless *et al.*, 2014], decentralized patrolling under limited communication [Alam *et al.*, 2015], and deceptive strategies that make environments appear more secure and improve resource use [Talmor and Agmon, 2017]. In these studies, the adversary is typically modeled abstractly

or algorithmically, and human decision-making is not directly observed or measured.

3 Experimental Setup

We used virtual reality to conduct an experiment that required participants to role-play as a school shooter. The experiment was conducted in a virtual replication of Columbine High School, the site of a well-documented school shooting in the United States. The environment was populated with non-player characters (NPCs) acting as students and teachers. For five minutes, the participant attempted to shoot NPCs, while two robots (one per floor) tried to intervene.

Participants. University students and faculty were recruited using poster and email advertisements. Because 96–97% of school shooters are male, participation was restricted to males [Peterson and Densley, 2023]. Additional inclusion criteria required participants to be at least 18 years old and not prone to motion sickness. Participants were initially paid \$15 per hour; this rate was later increased to \$30 per hour to improve recruitment. Individuals were not allowed to participate more than once. This study was IRB approved.

Independent Variables. As an initial investigation of the potential for robots to intervene and distract school shooters, we manipulated two independent variables:

- **Robot approach** was either *passive*, in which the robot followed the participant by five meters, or *aggressive*, where the robot raced to the shooter’s predicted future position (five seconds ahead).
- **Distraction level** increased from auditory cues only (low), to combined auditory and visual cues (medium), to the addition of robot-deployed smoke that partially obscured vision (high).



Figure 2: Participants were seated in a swivel chair with an HTC Vive headset. Shoe interfaces (Cybershoes) were strapped over their shoes allowing them to move through the simulation by making a walk motion. Valve Index controllers were strapped to hands for hand movements and shooting.




Figure 3: The simulation environment (left) was a replica of Columbine High School (right). A finite-state machine was used to control NPC behavior. The visible NPCs are in a relaxed state. The real image on the right was released by Bill Ockham in 2023 [Ockham, 2023].

Dependent Variables. The following dependent variables were recorded for each participant:

- *Number of victims (V)*: A count of the number of NPCs shot during the trial.
- *Robot gaze (G^R)*: A proportion of total frames where the participant’s eyes were looking directly at the robot.

The selected variables inform us about the potential value of the research and why our approach might work. All gaze and behavioral data were recorded at 2 Hz; brief eye closures were retained as natural behavior (blink rate: $M = 2.6\%$, $SD = 2.2\%$), with no systematic tracking loss across conditions. The raw data is available on GitHub. Due to the sensitive nature of the scenario, the simulation environment is available from the authors upon request.

Hypotheses. A total of 45 hypotheses were pre-registered on the Open Science Framework (OSF), many of which were exploratory. For this paper, we focus on a subset² of hypotheses related to the selected dependent variables (V , G^R). The full preregistration, including hypotheses and analyses not reported in detail here, is available on OSF.

- H1. The number of victims (V) will decrease with increased level of distraction.
- H2. The number of victims (V) will decrease with increasingly aggressive robot approach.
- H3. The shooter’s attention to the robot (G^R) will increase with increased distraction level.
- H4. The shooter’s attention to the robot (G^R) will increase with increasingly aggressive robot approach.
- H5. The presence of a robot will reduce the number of victims (V) relative to the baseline condition without a robot.

²Restricting analysis to a subset of hypotheses did not affect our use of multiple-comparison correction [Holm, 1979]. The total number of hypotheses within each family was still used to determine the appropriate adjustment.

Procedure. The procedure is depicted in Figure 4. Upon arrival to the lab, the participant was asked to complete a consent form. The participant was then outfitted with virtual reality equipment, enabling them to view a virtual environment that resembled the lab where the experiment was taking place. An NPC resembling the experimenter then led the participant first through a training procedure which allowed them to practice moving and shooting in the virtual environment. Next, the NPC led participants on a tour of the school highlighting points of interest. Because most school shooters have detailed knowledge of the school, the purpose of this tour was to increase the subject’s familiarity with the school. The participant was then given a virtual handgun. Finally, NPCs representing students, faculty, and staff arrived on campus. The shooting began once the participant fired their first shot and




Figure 4: Each participant followed the procedure above upon arrival. Once the experiment was complete, they were paid. Participants could quit at any time if they felt physical and/or mental discomfort and they would still be paid.

continued for a total of five minutes, reflecting the duration of the majority (approximately 69%) of documented active shooter incidents [Federal Bureau of Investigation, 2014]. In all test conditions, it was assumed that surveillance cameras could provide the robot with the shooter’s current location, which then served as input to our shooter position model to predict the future position of the shooter.

Non-Player Characters (NPCs). Non-player characters (NPCs) were controlled by a finite state machine (FSM), a common practice for modeling evacuees [Sheeba and Jayaparthi, 2021; Che *et al.*, 2015; Kielar *et al.*, 2014]. The FSM was designed to mimic the “Run. Hide. Fight” response recommended by the Department of Education [U.S. Department of Education, 2013]. Fighting is difficult to accurately simulate and rarely occurs during actual school shootings, so this option was disregarded. The resultant FSM is depicted in Figure 5. The default state of the NPC was relaxed. When triggered by seeing the shooter, victim, or alarmed individual, the NPC would transition to an alarmed state. From the alarmed state, the NPC would transition to a hiding state with a 75% probability or to an escaping state with 25% probability. From the hiding state, the NPC would transition to a running state. From the escaping state, the NPC would transition to a crouching state. From the running state, the NPC would transition to a crouching state. From the crouching state, the NPC would transition to a running state. From the crouching state, the NPC would transition to a running state.

Robots. The robots were designed to simulate the Agile-X Ranger Mini 2.0 (RM2) shown in Figure 6. Similar to the RM2, the simulated robot could move in all directions. The RM2 has a maximum speed of 1.7 m/s, which is slightly faster than the average walking speed [Fitzpatrick *et al.*, 2006]. Accordingly, we set the maximum speed of the robots to be 130% of the walking speed of the participant in the simulation. The robot’s *approach behavior* was driven by a multi-channel LSTM encoder-decoder model [Xue *et al.*, 2018; Pfeiffer *et al.*, 2018], which has been successfully applied to




Figure 5: A finite state machine is used to guide the behavior of each NPC. The column of ovals describes an NPC’s possible behavioral states for the experiment. The lines denote transitions between different states. The numbers associated with a transition represent the probability of the transition. NPCs remain in a ‘Relaxed’ state until they are triggered by seeing the shooter or alarmed NPC. The column of rectangles describes the animation used to generate each behavior in the simulation.



Figure 6: Comparison of the real robot (left) to the simulated robot (right) used in the experiment.

pedestrian trajectory prediction. Details of this model can be found in the supplementary material online. The model takes sequences of prior locations and the relative position of objects as input, while the output is a sequence of predicted future locations. Shooter position was estimated via simulated surveillance cameras and processed off-board; latency and prediction error did not vary across participants or conditions.

To create *distractions*, arms with LED lights were added to the robot to wave back and forth. This design is based on previous designs that have been used on evacuation robots [Robinette *et al.*, 2014]. A flashing light was also added to the top of the robot, a loud speaker to act as a siren, and a security smoke device, similar to what has been recently sold as a security device [Smokescreen, 2024]. The smoke device releases puffs of smoke at a frequency of 1.0 Hz, in order to disrupt the vision of the shooter.

Size of Study. A pre-experiment power analysis based on pilot data indicated that approximately 30 participants per condition were required to detect medium effects on victim count. A full 2×3 factorial design (approach type $\times$ distraction level) plus a no-robot control would have required ≥ 210 participants. Two conditions were removed based on pilot results showing that passive approach with low or medium distraction produced no measurable effect on outcomes ($\Delta V \approx 3$, $p = .50$; $\Delta S \approx 0.6$, $p = .95$), consistent with minimal robot salience (mean gaze $\approx 6\%$). Analyses of the retained conditions remain valid for the stated hypothesis tests, but should not be interpreted as estimating the full approach $\times$ distraction interaction or as ruling out smaller effects from the omitted passive low- and medium-distraction conditions. A total of 150 participants were recruited and assigned a condition using block randomization.

Ecological Validity. Admittedly, the data from our study is not based on the actions of real school shooters. Using school shooters as a study population is not possible because most school shooters are killed during their attack or are currently serving prison sentences. To maximize ecological validity, we matched our participant demographics to those re-

ported for school shooters. Prior research shows that 96% of school mass shooters are male with a mean age of 22.8 years (± 8.1) [Peterson and Densley, 2023]. Accordingly, our participants were all male with a mean age of 23.5 years. Prior work has compared the behavior of participants role-playing as school shooters in VR with that of actual school shooters, using measures such as number of shots fired, shot accuracy, and number of victims [McClurg and Wagner, 2025; McClurg and Wagner, 2026]. These studies found the two populations to be statistically equivalent on these key measures. Finally, participants also reported high task seriousness ($M = 4.35$, $SD = 0.86$) and immersion ($M = 3.82$, $SD = 0.91$) on a 5-point Likert measure. Taken together, these findings serve as evidence that ecologically valid data can be produced by asking subjects to act as active shooters, although more work will be necessary to fully validate this approach.

Although it may seem counterintuitive, the behavior of most school shooters is generally rational and shows little evidence of being shaped by complex emotions (e.g., fear), crowd dynamics, police presence, or situational factors [Grider *et al.*, 2000; Fox News, 2018; Fox, 2024; State of Connecticut, 2012]. Rather, school shooters tend to stroll through the environment at a measured (not hurried) pace, engage (rather than avoid) police, and linger in specific areas of the environment. School shooter behavior is typically characterized as deliberative and rational (from the perspective of maximizing harm) and not chaotic, confused, or unpredictable [Grider *et al.*, 2000; Fox News, 2018; Fox, 2024; State of Connecticut, 2012]. Moreover, historical data investigating the shooter’s state of mind during a shooting does not exist, in part because many shooters perish during the attack [Peterson and Densley, 2023]. Most data related to shooters captures factors that could serve as a motive or be used to predict who will become a shooter. On the other hand, school shooter behavior during the shooting has been documented as deliberative, calm and focused on generating victims [Burrows *et al.*, 2022; John Pirro, 2013; Ausdemore, 2015].

Our study intentionally focuses on shooter behavior and we assume that although unique psychological trauma may motivate the shooter, the actions taken by most shooters are in pursuit of their goal to maximize casualties. Essentially, they move through the environment searching for potential victims and attack them. This assumption is supported by numerous police reports and investigations [Burrows *et al.*, 2022; John Pirro, 2013; Ausdemore, 2015]. As such, our experimental procedure is face valid; we ask participants to shoot as many NPCs as possible, a task described to each participant three times before the shooting period began and confirmed with post-experiment responses indicating participants took the task seriously.

4 Experimental Results

Data was collected from October – December 2024 from a total of 172 participants. Data from participants who did not finish the simulation due to physical discomfort (13), emotional discomfort (7), or hardware error (2) were discarded. For reasons described in the previous section, the subjects

were all male with an average age of 23.5 years. The race distribution of subjects was 43.3% Caucasian, 38.3% Asian, 3.3% Black, and 15.0% other.

Hypothesis Testing. For statistical analysis, we conducted Welch’s unequal-variance *t*-tests with a standard significance level of $\alpha = 0.05$. Holm-Bonferroni correction [Holm, 1979] was applied to reduce the likelihood of Type I family-wise error, whereas families were determined by grouping strongly correlated dependent variables.³ Given approximate normality across conditions ($|skew| < 1$, $|kurtosis| < 1.1$) and $n \geq 30$ per condition, Welch’s *t*-tests were appropriate; confirmatory Generalized Linear Model (GLM) and negative binomial analyses yielded consistent conclusions and are reported in the supplementary material online. Unless explicitly stated, we report uncorrected significance values. For completeness, Table 1 presents *p*-values with and without correction.

Our first hypothesis (**H1**) is partially supported. Using an aggressive approach, the effect of using light distractors by the robot on the number of victims ($V_{low} \rightarrow V_{med}$) was not significant, $t(57.99) = 0.863$, $p = 0.392$. The effect of smoke distractors on the number of victims ($V_{med} \rightarrow V_{high}$), however, was significant, $t(57.97) = 3.146$, $p = 0.003$. Yet, the effect of both light and smoke distractors on the number of victims ($V_{low} \rightarrow V_{high}$) was significant, $t(57.99) = 4.039$, $p < .001$. Our second hypothesis (**H2**) is not supported. In the retained high-distraction comparison, aggressive approach did not significantly reduce the number of victims relative to passive approach ($V_{pass} \rightarrow V_{agg}$), $t(57.06) = 1.704$, $p = 0.094$. Our third hypothesis (**H3**) is partially supported. With an aggressive approach, the effect of light distractors on the shooter’s attention to the robot ($G_{low}^R \rightarrow G_{med}^R$) was significant, $t(46.98) = -2.477$, $p = 0.017$. In contrast, the effect of smoke distractors on the shooter’s attention ($G_{med}^R \rightarrow G_{high}^R$) was not significant, $t(54.78) = 1.906$, $p = 0.062$. The effect of lights and smoke on the shooter’s attention ($G_{low}^R \rightarrow G_{high}^R$) was not significant, $t(53.95) = -0.467$, $p = 0.642$. Our fourth hypothesis (**H4**) is partially supported. With a high level of robot distraction, the effect of the increasingly aggressive approach on the shooter’s attention ($G_{pass}^R \rightarrow G_{agg}^R$) was significant, $t(42.46) = -6.524$, $p < .001$. Our fifth hypothesis (**H5**) is supported. The effect of the robot presence on the number of victims ($V_{none} \rightarrow V_{robot}$) was significant, $t(40.38) = 2.782$, $p = 0.008$. More specifically, the best-case condition—approaching aggressively with smoke and lights—produced significantly fewer victims than having no robot, $t(54.71) = 4.359$, $p < .001$. Figure 7 presents the distributions of the primary dependent variables per condition, with significance annotated. Across participants, the number of shots directed at the robot was negatively correlated with total victims ($r = -.37$, $p < .001$), indicating that drawing fire toward the robot was associated with fewer casualties.

The effects of robot intervention are summarized in Table 1, where we report both original and adjusted *p*-values,

³There were strong correlations between victims and shots fired ($r = 0.69$), as well as gaze towards and shots fired at the robot ($r = 0.62$).




Figure 7: Victims (left) and gaze at the robot (right) across robot conditions. Error bars show means with 95% CI. Bold labels indicate differences from the no-robot control. Asterisks denote significance from Welch’s t -test: $*p < 0.05$, $**p < 0.01$, $***p < 0.001$.

descriptive statistics, and effect size. With respect to the number of victims (V), there is some evidence of a small-to-medium effect when using an aggressive (versus passive) robot approach. There was also some evidence of a small effect in using lights (versus no lights) as distractors. The effect of adding smoke, as well as adding lights and smoke, was large and statistically significant, reducing the number of victims by 35.6% and 41.3%, respectively. Finally, the best-case robot condition, approaching aggressively with smoke and lights, had a large and significant effect, reducing the number of victims by 46.6% as compared to the control (no robot) condition.

With respect to the shooter’s attention (G^R), there was a large and statistically significant effect when using an aggressive (versus passive) approach, which increased the shooter’s gaze toward the robot by 190.5%. There is also evidence

(significant before adjustment) of a moderate-to-large effect when adding lights as distractors, increasing the shooter’s gaze towards the robot by 35.1%. Adding smoke as a distractor had a moderately adverse effect, although this observation is more likely due to participants having greater difficulty tracking the robot in a smoky environment, rather than because they attended to the robot less.

Survey Responses. After completing the experiment, participants were asked to rate the importance, helpfulness, and perceived danger of the study. The majority rated the study as very to extremely important (86%) and helpful (75%), while most reported that the experience was not at all or only slightly dangerous (75%). These findings suggest that participants engaged seriously with the task and viewed the research as valuable, supporting the credibility of the data.

Hypothesis	Effect $A \rightarrow B$	p-value (original)	p-value (adjusted)	Mean (Condition A)	Mean (Condition B)	Mean Difference (%)	Effect Size, $ d $
1	$V_{low} \rightarrow V_{med}$	0.392	0.392	32.10 ± 12.57	29.23 ± 12.73	-8.9%	0.22
	$V_{med} \rightarrow V_{high}$	0.003 ✓	0.021 ✓	29.23 ± 12.73	18.83 ± 12.45	-35.6%	0.84
	$V_{low} \rightarrow V_{high}$	< .001 ✓	0.001 ✓	32.10 ± 12.57	18.83 ± 12.45	-41.3%	1.06
2	$V_{pass} \rightarrow V_{agg}$	0.094	0.469	24.80 ± 14.17	18.83 ± 12.45	-24.1%	0.45
3	$G_{low}^R \rightarrow G_{med}^R$	0.017 ✓	0.118	9.55 ± 3.69	12.90 ± 6.26	+35.1%	0.65
	$G_{med}^R \rightarrow G_{high}^R$	0.062	0.371	12.90 ± 6.26	10.08 ± 4.89	-21.9%	0.50
	$G_{low}^R \rightarrow G_{high}^R$	0.642	1.000	9.55 ± 3.69	10.08 ± 4.89	+5.5%	0.10
4	$G_{pass}^R \rightarrow G_{agg}^R$	< .001 ✓	< .001 ✓	3.47 ± 2.43	10.08 ± 4.89	+190.5%	1.71
5	$V_{none} \rightarrow V_{robot}$	0.008 ✓	0.057	35.23 ± 15.99	26.24 ± 13.93	-25.5%	0.60
	$V_{none} \rightarrow V_{best}$	< .001 ✓	0.001 ✓	35.23 ± 15.99	18.83 ± 12.45	-46.6%	1.14

Table 1: Summary of the reported effects of robot intervention. Post-hoc adjustment of p-values is determined by Holm-Bonferroni [Holm, 1979] method. Checkmarks indicate respective p-values are significant ($p < 0.05$).

5 Discussion

Our findings suggest that robot-based harm reduction may operate through two related mechanisms: capturing shooter attention and impairing target engagement. Aggressive approach increased gaze toward the robot, consistent with attention capture. However, the strongest reduction in victim counts occurred under high distraction, where smoke appeared to reduce visibility rather than increase sustained gaze toward the robot. Complementary analyses in the supplementary material online suggested that high-distraction conditions reduced firing behavior without substantially reducing movement, supporting the interpretation that smoke disrupted target engagement rather than simply slowing participants down. These findings have design implications for future response robots. First, multi-modal distractors should be prioritized over single-channel cues such as lights or sirens. Second, an aggressive approach strategy appears to enhance distraction by forcing visual confrontation, though future systems may need to balance aggressiveness against safety for bystanders. Third, our results motivate real-time adaptation: a robot might escalate from passive to aggressive behaviors or vary distractor intensity based on the shooter’s current state (e.g., proximity to victims, orientation toward exits). These implications should guide future testing rather than suggest deployment readiness.

Ethical Considerations. This research raises ethical issues related to both eventual deployment and participant safety. The risk of over-reliance, false positives and negatives, and intentional misuse (e.g., triggering the robot for fun) must be weighed against the system’s potential for saving lives. Although the robot’s distractors were selected to avoid direct physical force and reduce risk in false-positive scenarios involving students or staff, communication and rescue efforts could still be impacted if the robot malfunctions, gives students a false sense of escapability, or obstructs first responders. Because of the high-stakes nature of this work, additional research will be needed before deploying or relying on such a system. In particular, any deployed system would require evaluation with school communities, staff, administrators, and first responders, especially for design choices such as smoke in enclosed spaces.

In addition to societal and deployment risks, we considered the potential impact of this study on the participants themselves. We found no evidence that participants were emotionally harmed by participating in this research. Post-experiment surveys indicated that over 75% of participants felt the study was either not-at-all or only slightly dangerous. Participants were told that they could quit at any point if they felt physical and/or emotional discomfort and would still receive full payment; only 7 out of 172 participants chose to withdraw for this reason. To minimize risk further, we provided all participants with the contact information for university mental health services should they experience distress after leaving the site of the experiment. The experiment was intentionally designed to be face valid, avoiding deception or surprise, and was approved by the Institutional Review Board.

Limitations. Several limitations should be considered when interpreting these results. First, this study necessarily

relied on participants acting as shooters rather than real shooters, as direct access to actual school shooters is not feasible. Participants also knew they were acting within a simulation with no real victims or legal consequences, which may have affected their motivation, risk perception, and responses to robot distraction. However, prior work and our own participant surveys indicate that behavior in this paradigm is statistically equivalent to historical data on key measures such as shot rate, victim rate, and accuracy [McClurg and Wagner, 2025; McClurg and Wagner, 2026]. Second, the experiment used scripted NPC behaviors based on the widely recommended “Run. Hide. Fight” protocol, which reflects observed victim behavior but may not capture the full variability and social dynamics of real incidents [Ausdemore, 2015; John Pirro, 2013]. Third, robot behaviors were limited to a small set of distraction conditions and did not account for potential adversarial adaptation or multiple coordinated attackers. Instead, this study intentionally isolated the effect of several key distractors in the most common shooting paradigm: single-shooter incidents [Peterson and Densley, 2023]. Finally, the simulated robots were idealized and did not include sensor noise, actuation error, or real-world delays. Nonetheless, the distraction modalities themselves – loud noises, bright lights, and smoke – are physical phenomena likely to have similar effects in practice.

6 Conclusion

This paper explores the application of robots to delay and/or distract school shooters. Our human subject experiment, which asks participants to act as a school shooter in virtual reality, provides evidence that a robot can be used to reduce the number of victims in a shooting event. One specific robot condition – aggressively approaching the shooter with both lights and smoke – significantly reduced the number of victims by 46.6%, compared to having no robot. Beyond the empirical results, this work contributes a framework for safe, reproducible study of high-risk and adversarial human-robot interaction scenarios. By combining immersive virtual reality with well-defined robot behaviors, our approach enables systematic and ethically viable testing of robot strategies under conditions that would be infeasible or unsafe to reproduce in the real world.

Future work should increase the fidelity and scope of simulation while moving toward staged, controlled validation outside VR. This includes modeling sensor noise, actuation error, and latency to better approximate real-world robot performance. Additional experiments could also generate empirical models of victims, staff, and first responders, allowing future simulations to capture more of the broader response system. Failure modes should also be studied, including possible side effects of distraction, such as whether smoke slows victim evacuation or interferes with rescue operations. Beyond simulation, controlled exercises with first responders would be needed to evaluate how such robots affect response behavior, communication, and rescue operations. Together, these efforts would support more generalizable design guidelines and the responsible development of robots for school safety.

Acknowledgments

This material is based upon work supported by the National Science Foundation under Grant No. IIS-2045146. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation. Portions of this paper were edited and refined with assistance from generative AI. The authors reviewed and take responsibility for all content.

References

- [Addington, 2009] Lynn A. Addington. Cops and cameras: Public school security as a policy response to columbine. *American Behavioral Scientist*, 52(10):1426–1446, 2009.
- [Agmon *et al.*, 2011] Noa Agmon, Gal A. Kaminka, and Sarit Kraus. Multi-robot adversarial patrolling: facing a full-knowledge opponent. *Journal of Artificial Intelligence Research*, 42:887–916, 2011.
- [Alam *et al.*, 2015] Tauhidul Alam, Matthew Edwards, Leonardo Bobadilla, and Dylan Shell. Distributed multi-robot area patrolling in adversarial environments. In *International Workshop on Robotic Sensor Networks*, 2015.
- [Ausdemore, 2015] Steve E. Ausdemore. Eliminating the lost time interval of law enforcement to active shooter events in schools. Technical report, Naval Postgraduate School Monterey United States, 2015.
- [Bachman *et al.*, 2011] Ronet Bachman, Antonia Randolph, and Bethany L. Brown. Predicting perceptions of fear at school and going to and from school for african american and white students: The effects of school security measures. *Youth & Society*, 43(2):705–726, 2011.
- [Borum *et al.*, 2010] Randy Borum, Dewey G. Cornell, William Modzeleski, and Shane R. Jimerson. What can be done about school shootings? A review of the evidence. *Educational Researcher*, 39(1):27–37, 2010.
- [Burrows *et al.*, 2022] Dustin Burrows, Joe Moody, and Eva Guzman. Investigative committee on the robb elementary shooting: Interim report. Technical report, Texas House of Representatives, Austin, TX, 2022.
- [Che *et al.*, 2015] Xiaodong Che, Yu Niu, Bin Shui, Jianbo Fu, Guangzheng Fei, Prashant Goswami, and Yanci Zhang. A novel simulation framework based on information asymmetry to evaluate evacuation plan. *The Visual Computer*, 31:853–861, 2015.
- [Cramer and Hauser, 2021] Maria Cramer and Christine Hauser. Digidog, a robotic dog used by the police, stirs privacy concerns. *The New York Times*, 27, 2021.
- [Crawford and Burns, 2015] Charles Crawford and Ronald Burns. Preventing school violence: Assessing armed guardians, school policy, and context. *Policing: An International Journal of Police Strategies & Management*, 2015.
- [Elklit and Kurdahl, 2013] Ask Elklit and Sessel Kurdahl. The psychological reactions after witnessing a killing in public in a danish high school. *European Journal of Psychotraumatology*, 4(1):19826, 2013.
- [Federal Bureau of Investigation, 2014] Federal Bureau of Investigation. A study of active shooter incidents in the united states between 2000 and 2013, September 2014. Accessed 01-05-2026.
- [Fitzpatrick *et al.*, 2006] Kay Fitzpatrick, Marcus A. Brewer, and Shawn Turner. Another look at pedestrian walking speed. *Transportation research record*, 1982(1):21–29, 2006.
- [Fox News, 2018] Fox News. Chilling animation: Parkland shooter’s movements in school, 2018.
- [Fox, 2024] Melissa Fox. Critical incident review: Active shooter at robb elementary school. Critical incident review, VPC: Violence Policy Center, 2024.
- [Grider *et al.*, 2000] Kathi Grider, Rick L. Young, Kathleen Battan, and Jennifer Watson. Columbine high school shootings: Jefferson county sheriff’s office report. Report, Jefferson County Sheriff’s Office, Golden, CO 80401, 2000. Version 2000-05-15.
- [Holm, 1979] Sture Holm. A simple sequentially rejective multiple test procedure. *Scandinavian journal of statistics*, pages 65–70, 1979.
- [Huang *et al.*, 2019] Li Huang, MengChu Zhou, Kuangrong Hao, and Edwin Hou. A survey of multi-robot regular and adversarial patrolling. *IEEE/CAA Journal of Automatica Sinica*, 6(4):894–903, 2019.
- [James and McCallion, 2013] Nathan James and Gail McCallion. School resource officers: Law enforcement officers in schools, 2013.
- [John Pirro, 2013] John Pirro. Newtown police response to shooting under review. *News-Times*, 2013.
- [Jonson, 2017] Cheryl Lero Jonson. Preventing school shootings: The effectiveness of safety measures. *Victims & Offenders*, 12(6):956–973, 2017.
- [Kielar *et al.*, 2014] Peter M. Kielar, Oliver Handel, Daniel H. Biedermann, and André Borrmann. Concurrent hierarchical finite state machines for modeling pedestrian behavioral tendencies. *Transportation Research Procedia*, 2:576–584, 2014.
- [Knutsson, 2025] Kurt Knutsson. China rolls out its crime-fighting ball to chase down criminals. <https://www.foxnews.com/tech/china-rolls-out-its-crime-fighting-ball-chase-down-criminals>, 2025.
- [Lundberg and Christensen, 2007] Carl Lundberg and Henrik I. Christensen. Assessment of man-portable robots for law enforcement agencies. In *Proceedings of the 2007 Workshop on Performance Metrics for Intelligent Systems*, pages 76–83, 2007.
- [McClurg and Wagner, 2025] Christopher A. McClurg and Alan R. Wagner. Using virtual reality to simulate and study

- the movements of school shooters. In *2025 Annual Modeling and Simulation Conference (ANNSIM)*, pages 1–13. IEEE, 2025.
- [McClurg and Wagner, 2026] Christopher A. McClurg and Alan R. Wagner. Developing shooter movement models from virtual reality simulation. *Simulation: Transactions of the Society for Modeling and Simulation International*, 2026.
- [Mongan *et al.*, 2009] Philip Mongan, Schnavia Smith Hatcher, and Tina Maschi. Etiology of school shootings: Utilizing a purposive, non-impulsive model for social work practice. *Journal of Human Behavior in the Social Environment*, 19(5):635–645, 2009.
- [Nayyar and Wagner, 2022] Mollik Nayyar and Alan R. Wagner. Modeling evacuee behavior for robot-guided emergency evacuation. In *SEANAVBench Workshop at IEEE International Conference on Robotics and Automation (ICRA)*, 2022.
- [Nayyar and Wagner, 2023] Mollik Nayyar and Alan R. Wagner. Near real-time position tracking for robot-guided evacuation. In *SEANAVBench Workshop at IEEE International Conference on Robotics and Automation (ICRA)*, 2023.
- [Nguyen and Bott, 2001] Hoa G. Nguyen and John P. Bott. Robotics for law enforcement: Applications beyond explosive ordnance disposal. In *Enabling Technologies for Law Enforcement and Security*, volume 4232, pages 433–454. SPIE, 2001.
- [Ockham, 2023] Bill Ockham. Crime scene photo.chs library_klebold’s trenchcoat. <https://archive.org/details/crime-scene-photo-chs-library-klebolds-trenchcoat>, February 2023. JPEG image, Internet Archive.
- [Peterson and Densley, 2023] Jillian Peterson and James Densley. The violence project database of mass shootings in the united states (version 7). <https://www.theviolenceproject.org>, 2023. Accessed 12-29-2025.
- [Pfeiffer *et al.*, 2018] Mark Pfeiffer, Giuseppe Paolo, Hannes Sommer, Juan Nieto, Rol Siegwart, and Cesar Cadena. A data-driven model for interaction-aware pedestrian motion prediction in object cluttered environments. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5921–5928. IEEE, 2018.
- [Quigley *et al.*, 2009] Morgan Quigley, Brian Gerkey, Ken Conley, Josh Faust, Tully Foote, Jeremy Leibs, Eric Berger, Rob Wheeler, and Andrew Ng. Ros: an open-source robot operating system. In *ICRA Workshop on Open Source Software*, 2009.
- [Riedman, 2025] David Riedman. K–12 school shooting database. <https://k12ssdb.org>, 2025. Accessed 11-20-2025.
- [Robinette *et al.*, 2014] Paul Robinette, Alan R. Wagner, and Ayanna M. Howard. Assessment of robot guidance modalities conveying instructions to humans in emergency situations. In *The 23rd IEEE international symposium on robot and human interactive communication*, pages 1043–1049. IEEE, 2014.
- [Sheeba and Jayaparvathy, 2021] Angel A. Sheeba and Raman Jayaparvathy. Study of human factors contributing to fatal injury in a multi-floor building during emergency using finite state machines. In *2021 Asian Conference on Innovation in Technology (ASIANCON)*, pages 1–5. IEEE, 2021.
- [Sless *et al.*, 2014] Efrat Sless, Noa Agmon, and Sarit Kraus. Multi-robot adversarial patrolling: facing coordinated attacks. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*, pages 1093–1100, 2014.
- [Smokescreen, 2024] Concept Smokescreen. Security Smoke Products — Concept Smoke Screen. <https://www.smoke-screen.com/>, 2024. Accessed 04-12-2024.
- [State of Connecticut, 2012] State of Connecticut. Sandy hook elementary school shooting reports. State of Connecticut, 2012.
- [Suomalainen *et al.*, 2011] Laura Suomalainen, Henna Haravuori, Noora Berg, Olli Kiviruusu, and Mauri Marttunen. A controlled follow-up study of adolescents exposed to a school shooting—psychological consequences after four months. *European Psychiatry*, 26(8):490–497, 2011.
- [Talmor and Agmon, 2017] Noga Talmor and Noa Agmon. On the power and limitations of deception in multi-robot adversarial patrolling. In *IJCAI*, pages 430–436, 2017.
- [U.S. Department of Education, 2013] U.S. Department of Education. Guide for developing high-quality school emergency operations plans, 2013.
- [U.S. Department of Health and Human Services, 2015] U.S. Department of Health and Human Services. Results from the school health policies and practices study 2014. Atlanta, GA, 2015.
- [Xue *et al.*, 2018] Hao Xue, Du Q. Huynh, and Mark Reynolds. Ss-lstm: A hierarchical lstm model for pedestrian trajectory prediction. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1186–1194. IEEE, 2018.