Security Journal
https://doi.org/10.1057/s41284-024-00423-7

ORIGINAL ARTICLE

Bridging the gap between criminology and computer vision: A multidisciplinary approach to curb gun violence

Tyler E. Houser[1] · Alan McMillan[2] · Beidi Dong[3]

Accepted: 15 February 2024

Abstract

Gun violence significantly threatens tens of thousands of people annually in the United States. This paper proposes a multidisciplinary approach to address this issue. Specifically, we bridge the gap between criminology and computer vision by exploring the applicability of firearm object detection algorithms to the criminal justice system. By situating firearm object detection algorithms in situational crime prevention, we outline how they could enhance the current use of closed-circuit television (CCTV) systems to mitigate gun violence. We elucidate our approach to training a firearm object detection algorithm and describe why its results are meaningful to scholars beyond the realm of computer vision. Lastly, we discuss limitations associated with object detection algorithms and why they are valuable to criminal justice practices.

Keywords Gun violence · CCTV · Situational crime prevention · Deep learning · Automated firearm detection

Introduction

Gun violence is a significant and growing problem in the United States. In 2021, firearm-related homicides accounted for approximately 20,966 deaths in the United States (Simon et al. 2022). The firearm-related homicide estimates represent an 8.3% increase from 2020, which was a 34% increase from 2019 and a 75% increase over the last decade (CDC 2022). Furthermore, gun violence disproportionately

✉ Tyler E. Houser
thouser@gmu.edu

[1] Department of Criminology, Law and Society, George Mason University, Enterprise Hall, 4400 University Drive, MS 4F4, Fairfax, VA 22030, USA

[2] Department of Radiology, School of Medicine and Public Health, University of Wisconsin, 1111 Highland Avenue, Madison, WI 53705, USA

[3] Department of Criminology, Law and Society, George Mason University, 354 Enterprise Hall, 4400 University Drive, MS 4F4, Fairfax, VA 22030, USA

Published online: 02 April 2024

affects specific groups in the United States. African Americans experienced an increase of roughly 39% in firearm-related homicides from 2019 to 2020 and experienced the highest firearm homicide rate by race/ethnicity in 2021 (CDC 2022; Simon et al. 2022). Additionally, the second leading cause of death among individuals aged 12–19 is now homicide, in which firearms account for 87% of cases (Kolbe 2020). Lastly, active shooter and mass shooting events are also rising in the United States. The FBI found that only three active shooter events occurred in 2000, but 61 occurred in 2021, representing a 1933.33% increase (Blaire and Schweit 2014; Federal Bureau of Investigation 2022).

Criminology and other social science scholars have dominated the literature in proposing gun violence solutions and implications in the past (e.g., Rosenfeld et al. 2014; Sherman and Rogan, 1995; Makarios and Pratt 2012; Braga et al. 2008; Koper et al. 2015). More recently, given the considerable threat to public safety and public health in the United States, other disciplines have begun to offer solutions to prevent gun violence, most notably computer vision scholars. In the past five years, a growing body of literature has focused on using artificial intelligence to curb gun violence. Computer vision scholars have used artificial intelligence to train algorithms to detect firearms in real time, which can be implemented in closed-circuit television (CCTV) systems (Ashraf et al., 2022; Narejo et al. 2021; Garza and Vega 2021; Bhatti et al. 2021; Ahmed et al. 2022).

Despite the application of such algorithms concerning the criminal justice system, to our knowledge, very few criminology scholars have explored object detection algorithms to mitigate gun violence (e.g., Idrees et al. 2018). Not only does this leave a significant gap in the literature in most social science journals, but it also reduces the likelihood of optimal implementation in real-world settings. Although computer vision scholars are more formally trained in these artificial intelligence techniques, criminologists can guide how, when, and where implementations will be most effective. Thus, the disciplines must collaborate to optimally implement state-of-the-art security technologies to enhance public safety.

Therefore, in this work, we demonstrate how artificial intelligence can assist the criminal justice system in preventing and mitigating the consequences of gun violence. We provide a conceptual overview of a multidisciplinary approach integrating criminological frameworks and computer vision techniques to help curb the gun violence epidemic in the United States. Specifically, we put forth a theoretical framework, situational crime prevention, that explains how implementing firearm detection algorithms can improve current CCTV limitations related to gun violence. We also provide a general description of object detection and discuss relevant work on automated firearm detection. We then demonstrate our data collection methods and how we trained a firearm object detection algorithm. Additionally, we explain why our algorithm's results are meaningful to the criminal justice system and why they should transfer to real-world settings. Lastly, we describe limitations associated with detection algorithms and future research directions.

Theoretical justification for automated firearm detection algorithms

Situational crime prevention and CCTV

Although research on artificial intelligence-based firearm detection algorithms is vastly growing, little empirical testing has been done to examine firearm detection algorithms in the real world. The state of the literature is largely conceptual. Additionally, to our knowledge, the literature has not incorporated a theoretical framework to explain *why* firearm object detection algorithms would reduce and prevent gun violence in the real world, limiting confidence that the algorithms would be successful. Although this work does not add empirical results (e.g., showing the success or failure of real-life applications of the algorithms), it does provide a theoretical mechanism, situational crime prevention, to explain *why* firearm object detection algorithms should prevent and mitigate gun violence in real-world settings.

Many criminology perspectives focus on distal issues that lead to crime, such as childhood factors, strain, or peers (Gottfredson and Hirschi 2019; Agnew 1992; Akers 1999; Snipes et al. 2019). However, when attempting to prevent crime in the present or immediate future, perspectives focused on distal factors do not provide a theoretical basis for crime prevention. Unlike most criminology theories, situational crime prevention focuses on the immediate environment where crime will occur (Clarke 2017). Essentially, situational crime prevention seeks to alter a vulnerable environment to reduce the number of criminal opportunities available or increase the perceived risks of apprehension for offenders, preventing or reducing crime.

Although various situational crime prevention strategies exist (Clarke 2017), we focus on one technique: closed-circuit television (CCTV). CCTV is a surveillance technology often used by law enforcement organizations (as well as other public and private entities). CCTV is "a system in which a number of video cameras are connected in a closed circuit or loop, with the images produced being sent to a central television monitor or are recorded" via a wireless, remote network (Goold 2004, p. 12; Ratcliffe and Rosenthal 2021). The use of CCTV is vastly growing across the United States. Nearly all law enforcement agencies that serve large populations use CCTV for various crime-reduction interventions (Reaves 2015).

Relevant to this work, CCTV reduces crime by strengthening the formal surveillance of potential crime settings, which increases the likelihood of offender apprehension and allows law enforcement (or other place managers) to respond to incidents more efficiently and effectively (Clarke 2017; Piza et al. 2019). However, the existing literature suggests that CCTVs do not reduce violent crime. In a meta-analysis and systematic review of 80 evaluations of CCTV, only drug, property, and vehicle crimes were found to be meaningfully reduced by CCTV interventions (Piza et al. 2019). Nonetheless, there is reason to believe that artificial intelligence can augment current CCTV practices (Skogan 2019). Specifically, CCTV systems embedded with automated firearm detection algorithms should help combat the gun violence epidemic through several mechanisms.

First, automated firearm detection algorithms create active CCTV monitoring with or without a human present. In their 80-study meta-analysis and systematic review, Piza et al. (2019) found that actively monitored systems are significantly more likely to prevent crime than passive systems. Active systems require live-feed CCTV footage to be monitored at all times, whereas passive systems only record and store footage. Relevant to this work, a firearm will not be detected in real time if a system is not actively monitored. Thus, by embedding a CCTV system with real-time detection powered by a computer vision algorithm, the system will be actively monitored for firearm-related incidents at all times, increasing the efficacy of CCTV in preventing gun crimes.

Second, actively monitored systems by humans only sometimes allow for real-time detection. Idrees et al. (2018) illustrate that human operators generally have numerous camera systems to monitor at any given time. Thus, even when a firearm is on live-feed CCTV (and the system is actively monitored), a human operator may miss it as they could be distracted by other footage or simply did not notice it (Idrees et al. 2018; Ratcliffe and Rosenthal 2021). Although object detection algorithms are not perfect (discussed in the limitations section), they can still augment current practices to reduce human errors and improve detection capabilities. By situating a firearm detection algorithm within a CCTV system, human operators are provided with a “second pair of eyes,” and the likelihood of detecting a firearm is significantly increased, helping to mitigate gun violence (Rigano 2019, p. 3).

Third, ordinary CCTV systems are typically not tailored to prevent firearm-related crime, which is why most crime-reduction effects are seen in property crimes (Piza et al. 2019). A core tenet of situational crime prevention is that an intervention must be specific to the desired crime type (Clarke 2017). Thus, an intervention focused on property crime might not apply to gun crime. However, if practitioners tailor CCTV systems to gun violence using firearm detection algorithms, it is more likely that a gun will be detected, and the intended crime will be prevented.

Lastly, earlier firearm detection will allow police to respond more quickly to a potentially dangerous incident. Although evaluations of situational crime prevention demonstrate the theory’s effectiveness in combating varying crime problems (Guerette 2009; Guerette and Bowers 2009; Clarke 2017), like all criminology interventions, situational crime prevention cannot prevent every crime from occurring. Nonetheless, even if a situational crime prevention intervention does not prevent its targeted crime, it can still be beneficial. Freilich et al. (2020) argue that situational crime prevention can reduce the total number of casualties stemming from public mass violence through varying mechanisms; most notable to this study is *improved response times*. By timely notifying law enforcement of a firearm-related incident, police can quickly arrive at the scene and neutralize the threat, limiting the number of casualties. Furthermore, a timelier notification will also reduce the response time of emergency medical services and result in quicker transport to trauma care, which might help to save the lives of wounded victims (Hatten and Wolff 2020). Thus, immediately detecting a firearm in live-feed CCTV footage is critical to responding to and mitigating gun violence efficiently and effectively.

Ultimately, situational crime prevention and relevant CCTV literature provide automated firearm detection algorithms with a theoretical framework to help explain a causal mechanism for gun violence reduction and mitigation. Automated firearm detection technologies are tailored explicitly to gun violence and increase formal surveillance of potential crime settings (Clarke 2017). Additionally, automated firearm detection algorithms should increase the likelihood that firearms are detected in real-time footage by assisting human operators and creating an actively monitored system (Idrees et al. 2018). Lastly, automated firearm detection algorithms might not prevent all shootings, but they can assist law enforcement and emergency medical services with responding to incidents more efficiently and effectively (Freilich et al. 2020). Although the present research is conceptual, the current theoretical framework provides confidence that firearm detection algorithms would succeed in real-world applications.

Object detection

Deep learning is a form of machine learning that uses neural networks to discover patterns in visual data by examining large amounts of data (LeCune et al. 2015). A subset of deep learning includes object detection, which typically uses convolutional neural networks (CNN) to analyze visual data (e.g., photographs). CNNs allow object detection algorithms to learn and recognize patterns of data, including classification and localization of objects of interest within images or videos. Although various object detection algorithms are publicly available, we use the You Only Look Once (YOLO) algorithm. YOLO is a state-of-the-art object detection algorithm due to its speed and accuracy of detection, making it effective and practical for automated firearm detection tasks (Wang et al. 2022). Because YOLO has been found to be state-of-the-art and is employed in this work, the next portion of the literature review focuses on experiments that used a YOLO algorithm to detect firearms in visual data.

Ashraf et al. (2022) trained their experiment using a YOLO-v5s approach and a custom CNN to compare the performance of the two for firearm detection in images and videos. The authors trained their approach on a 15,873-image dataset (focusing solely on handguns). The authors' YOLO-v5s approach yielded a precision and recall of 99% and 81% on images and a precision and recall of 94% and 93% on videos, outperforming the other tested CNN. Furthermore, YOLO's detection speed was 19 times faster than their custom CNN, providing evidence that it is suitable to detect firearms quickly in live-feed CCTV footage.

Narejo et al. (2021) trained their experiment using the YOLOv3 algorithm and compared it to three other object detection algorithms available for public use (i.e., Alexnet + SVM, Faster RCNN, CNN VGG-16). The authors trained their approach on an original dataset derived from open sources. Their results demonstrated that YOLOv3 outperformed the other algorithms with an accuracy of 98.89%, providing confidence in its reliability in detecting firearms in CCTV applications. The authors also note that YOLOv3 was computationally less expensive than other algorithms.

Like Narejo et al. (2021), Garza and Vega (2021) used YOLOv3 to detect firearms in video frames. The authors' dataset included 18,000 (augmented) images from CCTV camera video frames containing firearms. The authors' approach reached a precision of 85%, a recall of 81%, an $F1$-score of 83%, and a mean average precision of 85% at the 0.5 IoU threshold. Furthermore, the algorithm detected firearms in varying conditions, including worsened image qualities and occluded gun positions. Additionally, in another experiment, Rosales et al. (2021) also illustrate that the YOLO algorithm can detect armed individuals in various settings (i.e., low light, different backgrounds, varying image qualities) with a relatively high degree of precision and recall, making it suitable for CCTV systems that monitor unfavorable environments.

Bhatti et al. (2021) trained varying object detection algorithms (including YOLO) on multiple custom datasets to determine the best-performing approach. The datasets ranged in size from 1732 to 8327 images. The authors' testing revealed that the YOLOv4 algorithm had the highest performance (compared to the other publicly available detectors) with a mean average precision and $F1$-score of 91.73% and 91%. The YOLOv4 algorithm also yielded a 99% detection confidence in most cases and provided the fewest false positives and negatives, illustrating its ability to perform within real-world security settings. Lastly, Ahmed et al. (2022) trained a scaled YOLOv4 algorithm on a roughly 8000-image dataset derived from various open sources. The authors' approach resulted in a 92.1% mean average precision. The authors also found success using high-performance and low-cost GPUs (similar to Narejo et al. 2021), providing confidence that YOLO is suitable for resource-depleted organizations.

The aforementioned studies demonstrate the effectiveness of object detection algorithms, specifically YOLO. Most work exceeded a precision of 90%, and all work exceeded a recall of 80%. Thus, the algorithms correctly identified and detected firearms in most instances throughout the various experiments. Although the algorithms were not perfect in any study, they still provide an advantage to current uses of CCTV. Implementing object detection algorithms to assist human operators will help increase the likelihood of real-time firearm detection, especially considering CCTV's previously outlined limitations. One must also consider that the capabilities of object detectors will only improve in the future, furthering the case for their application in the criminal justice system. Taken together, the current literature suggests that YOLO is a state-of-the-art object detection algorithm and should be reliable for real-time firearm detection in real-world applications.

Training a firearm detection algorithm

Algorithm: YOLO

Computer vision scholars and practitioners have released several ready-to-use detection algorithms for public consumption (e.g., Wang et al. 2022). The ready-to-use algorithms provide criminologists and social science scholars opportunities to train automated firearm detectors and produce state-of-the-art results (without being experienced coders or computer scientists). The following sections demonstrate a ready-to-use object detection algorithm that we trained to detect various firearms. The experiment should help illustrate to scholars outside of the computer vision field how such algorithms function and what their results mean for public safety and security applications.

We leveraged the YOLOv7x CNN to train our firearm object detection algorithm. Wang et al. (2022) developed and implemented YOLOv7 for public use. As previously mentioned, YOLO is a state-of-the-art object detection algorithm that stands for: You Only Look Once (Wang et al. 2022). YOLO detects objects using multiple techniques that are reduced to a single computation. The algorithm divides an image into grid cells to detect any objects of interest (Redmon et al. 2015). If the algorithm detects an object of interest in any grid cells, the algorithm calculates a predicted bounding box (detailing where an object is in the image) and locates the center of the object of interest. After predicting where an object of interest is in the image, the algorithm calculates an intersection over union (IoU) score to determine whether the labeled bounding box (ground truth data) overlaps with the predicted bounding box. This score explains how well the algorithm predicted where an object was in an image. In addition to the IoU score, the algorithm provides a confidence score of the class probability. Thus, how confident is the algorithm that a detected object is the predicted class (e.g., is the object a firearm)? Because YOLO completes the previous steps simultaneously, it provides an increased speed in detection, making it practical for real-time weapon detection in CCTV systems.

Data

Our firearm detection algorithm is trained on still images. Despite training the algorithm on still images (like most prior work), the algorithm can then be used for real-time video detection. Thus, the trained algorithm applies to CCTV technologies. Furthermore, the algorithm might see an increase in performance when exposed to video data (compared to still images used in this experiment). Rather than being assessed on single still frames, live-feed video provides numerous frames per second. Given YOLO's high speed of detection and ability to handle large amounts of data (i.e., numerous video frames), the additional opportunities to identify an object as a firearm might reduce false negatives (or missed detections) and increase overall efficacy (Ashraf et al., 2022).

We derived the bulk of our data from publicly available data sources. Although many firearm object detection datasets exist, most are not publicly available or do not contain appropriate labeling. Nonetheless, we located four firearm object detection datasets to help satisfy the requirements for training. Gu et al.'s (2022) dataset accounts for slightly less than half of our data, which includes 5000 challenging images and labels of people and firearms with rich background features (i.e., real-world data). We also used data from Kaya et al. (2021), Qi et al. (2021), and Duran-Vega et al. (2021), which included various images of different firearms. In addition to publicly available datasets, we supplemented our data with images from surveillance footage (via Google Images), publicly available photographs on social media and news media, and movies and TV shows.

Although the publicly available datasets already included labels (i.e., where a firearm or person is located in an image), the data from other open sources did not. Thus, we had to label the open-source data (i.e., place bounding boxes around objects of interest) to fit the requirements of the YOLO algorithm (and create ground truth data). We fed the unlabeled data into a YOLOv5 algorithm already trained to detect firearms and people. The algorithm detected objects of interest in each image, providing labeled data as output. Thus, instead of manually annotating each firearm or person by hand, the algorithm labeled the data to meet YOLO's requirements. Each labeled prediction was manually analyzed to ensure accuracy and precision in detecting and labeling a firearm or person. Accurate predictions were added to the final dataset to train our final approach using YOLOv7x. The final dataset included over 11,000 images divided into training, validation, and testing sets (approximately 80%, 10%, and 10%).

Evaluation

We evaluated our approach by examining its precision, recall, $F1$-score, average precision, and mean average precision. Precision assesses the ability to make positive predictions (Shah 2022), that is, how many predictions were true positives out of the total number of positive predictions made. Recall assesses how many true positives were detected out of the total number of objects (i.e., firearms, persons) in a dataset. The $F1$-score is a weighted average between the precision and recall metrics, ranging from 0 (lowest accuracy) to 1 (highest accuracy).

Average precision is the area under the precision-recall curve (Shah 2022). A larger area under the precision–recall curve indicates a higher average precision. Essentially, higher precision and recall scores result in a greater average precision. Each class has a corresponding average precision. Mean average precision is the average of the average precision of all classes. Thus, the mean average precision considers the average precisions of the gun and person classes and produces a single score. We evaluate average precision and mean average precision at the 0.5 IoU threshold. Again, IoU examines the overlap between the ground truth bounding box (object labels) and the bounding box predicted by the algorithm. An IoU score greater than 0.5 is a positive match (or true positive), and anything less than 0.5 results in a false positive. The mean average precision provides insight into the approach's overall ability regarding classification (i.e., predicting a gun or person) and localization (i.e., predicted bounding box relative to ground truth labels). The higher the mean average precision, the more accurate the model detects firearms and people.

Experiment

We conducted all experiments using Google Colaboratory Pro in 2022. We used the YOLOv7x algorithm and initial weights to train our firearm detection algorithm (https://github.com/WongKinYiu/yolov7) (Wang et al. 2022). We set our experiment parameters to 16-batch size, 50 epochs, and 640 image size. We also applied YOLO's ready-to-use *p6* hyperparameter settings, which were observed to improve model performance. We trained our approach on two classes, person and gun, and most training images contained at least one of those classes. We trained the algorithm on varying firearms, including handguns, rifles, and shotguns, all mapping to the single class of gun. Although the focus of this experiment and its implications is the gun class, the included person class provides importance to training. The person class helps to provide context of how firearms are used in gun violence events (and who uses them). Therefore, the inclusion of the person class provided additional information to the algorithm when learning to identify firearms in the data (Gu et al. 2022).

Results

The overall precision of our approach was 96%, the precision for the person class was 97.3%, and the precision for the gun class was 94.6%. The overall recall was 90.2%, the person class's recall was 94.7%, and the gun class's recall was 85.6%. The overall $F1$-score was 93%, the person class's $F1$-score was 96%, and the gun class's $F1$-score was 89.9%. The average precision of the person and gun classes were 97.4% and 92%. Lastly, the mean average precision of the gun and person classes was 94.7%. Table 1 illustrates the experiment's results.

The figures below illustrate how the algorithm detects an object of interest for scholars outside the computer vision field. The figures are stills from security camera footage of school shooting incidents. Specifically, the algorithm successfully detected the firearms used by one of the Columbine shooters, the Parkland High School shooter, and the Uvalde Elementary School shooter. Although we trained our approach to detect persons and guns, the examples below only display detections for the gun class to highlight its ability to perform within a security setting. Figures 1, 2, and 3 illustrate successful detection examples.

Discussion of results and applicability to the criminal justice system

The gun violence epidemic is a significant and growing threat to the safety of the United States. Not only is ordinary or general gun violence on the rise in the United States, but mass shootings are also rising with each passing year (Simon et al. 2022; Blaire and Schweit 2014; Federal Bureau of Investigation 2022). Despite the dire risk to tens of thousands yearly in the United States, much gun violence research lacks multidisciplinary implications and solutions to solve the problem. Scholars and practitioners from various fields must combine their efforts to create a holistic response to gun violence. We set out to bridge the gap between computer vision and criminology. In doing so, we present a theoretical

Table 1 Experiment results

Class	P (%)	R (%)	$F1$ (%)	AP@0.5 (%)	mAP@0.5 (%)
All	96.0	90.2	93.0	–	94.7
Person	97.3	94.7	96.0	97.4	–
Gun	94.6	85.6	89.9	92.0	–

Fig. 1 Columbine High School Shooting, 1999

Fig. 2 Parkland High School Shooting, 2018

framework, situational crime prevention, and provide an overview of object detection algorithms to criminologists and other social science scholars.

The various metrics in the above experiment provide important implications for criminal justice. For instance, the precision of the firearm class reached 94.6%, demonstrating that our approach can effectively identify firearms. If an algorithm can identify objects of interest with higher precision, the algorithm will produce fewer false positives (Shah 2022; Ahmed et al. 2022), which is critical for implementations in the criminal justice system. If a detection algorithm has low precision, it will more often incorrectly identify non-firearm objects as firearms, possibly sending law enforcement to a location deemed 'dangerous'

Fig. 3 Uvalde Elementary School Shooting, 2022

when it is not. By sending police to a dangerous event (that is not actually dangerous), there is an increased risk that law enforcement officers will use force on undeserving people.

The recall of the firearm class neared 85%. Thus, our approach correctly detected when a firearm was present 85% of the time during the experiment. As Idrees et al. (2018) note, human operators often face challenges in monitoring live-feed CCTV footage effectively. The ability to supplement human operators with a "second pair of eyes" will inevitably increase the likelihood that a firearm is detected (Rigano 2019, p. 3; Idrees et al. 2018). Thus, an 85% recall rate will assist human operators in detecting when a firearm is present in live-feed footage (in most cases). Additionally, as previously mentioned, exposing the algorithm to live-feed footage (versus single photographs) might also result in a higher recall rate as there are more opportunities to detect an object as a firearm, increasing its applicability to the criminal justice system.

Furthermore, our approach yielded a relatively high $F1$-score. Again, the $F1$-score indicates the weighted mean between the precision and recall metrics (Shah 2022). An $F1$-score of nearly 90% reiterates that this approach can successfully identify firearms and support human operators in detecting guns in overwhelming live-feed footage.

Lastly, the average precisions of the gun and person classes were 92% and 97.4%, yielding a mean average precision of 94.7% at the 0.5 IoU threshold. A mean average precision of nearly 95% indicates high accuracy in localization (i.e., a predicted bounded box relative to the ground truth label) and classification (i.e., predicting that an object is a gun). Thus, the model can successfully predict whether an object is a firearm and accurately determine where a gun is located within visual data (in the vast majority of cases), furthering its applicability to CCTV interventions.

Our approach also performed well compared to the existing YOLO firearm detection literature. For instance, our gun class's precision (94.6%), recall (85%), and $F1$-score (89.9%) exceeded that of Garza and Vega's (2021) YOLOv3 approach, which achieved a precision of 85%, a recall of 81%, and a $F1$-score of 83%. Additionally, our recall metric for the gun class exceeded that of Ashraf et al.'s (2022) experiment (84.6%), but our precision metric for the gun class did not reach theirs (99%). It is worth noting that Ashraf et al.'s (2022) approach employed images solely focused on handguns and images derived from movies, which might result in higher precision than the real-world images of various firearms used in our experiment. Our approach's mean average precision was similar to Ahmed et al.'s (2021) scaled YOLOv4 algorithm (94.7% and 92.1%). However, Ahmed et al.'s (2021) experiment trained a firearm class and a non-firearm class (where our classes were firearm and person), potentially differentiating the mean average precisions. Taken together, this approach (compared with the existing literature) demonstrates that scholars outside the computer vision field can produce state-of-the-art firearm detection algorithms that can increase the efficacy of CCTV technologies.

Despite the many successful experiments of firearm detection algorithms on various data, it is still unknown of their true efficacy in the real world (as there is no empirical test). However, again, situational crime prevention provides confidence that the algorithms will be successful through several mechanisms. First, CCTVs embedded with firearm detection algorithms are tailored explicitly to gun violence compared to CCTVs without detection capabilities, increasing the likelihood of success (Clarke 2017). Second, implementing detection algorithms improves current CCTV capabilities to provide formal surveillance of potentially violent settings by creating continuously, actively monitored systems, which is critical for crime reduction (Piza et al. 2015, 2019). Furthermore, fewer human errors will occur with the assistance of detection algorithms, which increases the level of formal surveillance of crime settings and the number of potential firearm detections (Idrees et al. 2018). Lastly, the use of firearm detection algorithms will result in the earlier detection of a firearm. Early firearm detection will dispatch police (and other emergency services) to a scene more quickly, mitigating the consequences of a gun violence event (Freilich et al. 2020). Ultimately, CCTVs embedded with firearm detection algorithms should reduce or mitigate gun violence in real-world applications through varying mechanisms provided by situational crime prevention.

Other applications of object detection algorithms to situational crime prevention

Although this paper solely focuses on gun violence, it must be mentioned that scholars can situate object detection algorithms in situational crime prevention to combat other criminal justice problems. For instance, some work has trained detection algorithms to detect other dangerous objects besides firearms (i.e., knives and other bladed weapons) (Buckchash and Raman 2017; Kibria and Hasan 2017; Berardini et al. 2023). Two arenas that might specifically benefit from knife and blade detectors are schools and correctional facilities.

Although active shooter events in schools have increased over the past couple of decades in the United States (Katsiyannis et al. 2023), schools must also guard against bladed violence, which accounts for a relatively large portion of school violence in America (National Threat Assessment Center 2019). Additionally, state and federal prisons have experienced increased violence, especially homicide, in the past two decades (Carson 2021). Although it is unlikely that inmates can obtain a firearm while incarcerated, they might be more likely to make (or potentially acquire) a knife or other bladed weapons (e.g., Seiden 2022). Thus, not only should algorithms be trained to detect firearms (most notably in schools), but they should also be able to detect knives and bladed weapons to prevent other types of violence.

Knife and blade detection algorithms will help to prevent violence in schools and correctional facilities through the exact mechanisms presented for firearm detectors. First, embedding knife and blade detection algorithms in CCTV systems tailors the intervention to the type of violence seen in schools and correctional facilities (Clarke 2017). Second, using detectors in schools and correctional facilities provides a "second pair of eyes" when monitoring live-feed CCTV (Rigano 2019, p. 3), helping to strengthen formal surveillance. Lastly, increasing formal surveillance should reduce police (and teachers/administrators) and correctional officers' response times when responding to relevant incidents at schools and correctional facilities, limiting harm to potential victims.

Practitioners can also situate object detectors in situational crime prevention to mitigate other crime problems. For instance, research has associated varying pre-incident behaviors and precursors with specific crime types (Piza and Sytsma 2015; Idrees et al. 2018; Piza et al. 2019; Skogan 2019). Although it might be more resource-intensive and challenging to train approaches to detect specific behaviors (e.g., hand-to-hand drug transactions in open-air drug markets, "car-hopping"), doing so would increase formal surveillance of varying crime settings (Piza and Systma 2015; Idrees et al. 2018, p. 297). Again, by increasing formal surveillance, specifically tailoring CCTV to crime types, and reducing human errors, law enforcement can respond to varying crime problems more efficiently and effectively (Clarke 2017).

Furthermore, some work has applied object detection algorithms to TSA's use of baggage scanners (e.g., Liang et al. 2019; Sigman et al. 2020; McKay et al. 2022). Like CCTV systems, algorithms can be implemented in X-ray technologies to detect when dangerous objects are present in luggage. McKay et al. (2022) argue that object detection algorithms might improve TSA's efficiency and overall performance (regarding safety) during baggage screening. However, like our argument regarding firearm detectors, McKay et al. (2022) suggest that object detection algorithms should not entirely supplant current practices. Thus, object detection algorithms offer TSA officers a "second pair of eyes" when screening passengers (Rigano 2019, p. 3). Regarding situational crime prevention, object detection algorithms in TSA screening technologies would increase the risks for offenders (through target hardening TSA checkpoints) and increase the surveillance of all luggage, preventing a potential crime from occurring (Clarke 2017). Although these are just a few examples, they further show the relevance of object detection algorithms to situational crime prevention and why criminologists should leverage the artificial intelligence technique to combat various types of crime.

Limitations and future research directions

Although object detection algorithms are vastly improving, they have limitations. First, some object detection algorithms are computationally expensive. Although we trained our approach using the YOLOv7x algorithm, we did not have enough computing power to explore larger YOLOv7 models, which might have improved our results. Thus, future work should continue to create computationally affordable algorithms for all scholars and practitioners. Furthermore, scholars and practitioners with more available computational power should release their trained algorithms for public use.

Second, object detection algorithms suffer from false positives and negatives (Olmos et al. 2017). The algorithms may detect when an object of interest is not present or fail to detect when an object of interest is present. Our algorithm produced false positives and false negatives in some instances where firearms were occluded, when firearms were further away from the camera, or when image quality was poor (despite Garza and Vega 2021 and Rosales et al. 2021 finding success with lesser quality images). Other instances of false positives include the detection of objects that resemble firearms (e.g., phones). Future research must continue to mitigate false-positive and false-negative risks associated with firearm detection algorithms to ensure their effectiveness in real-world applications. For those scholars unfamiliar with computer vision, we illustrate an example of a false negative in Fig. 4 and examples of false positives (along with a correct detection) in Fig. 5.

Although algorithmic tweaks and advances will help reduce issues associated with false positives and negatives, scholars and practitioners must also focus their efforts on data. An object detection algorithm is only as good as the data available

Fig. 4 False negative

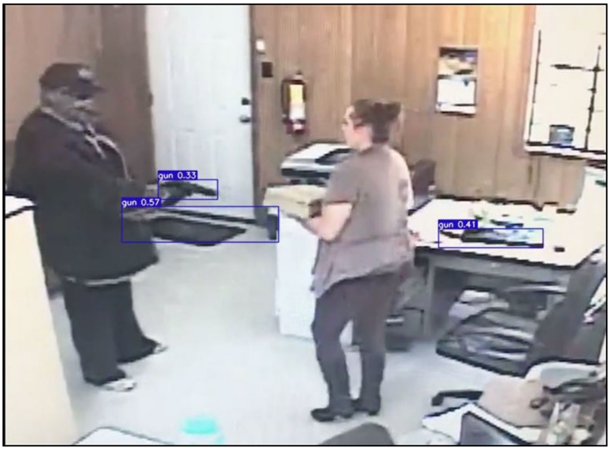

Fig. 5 False positive

for training (McDonald 2022). Thus, researchers and practitioners must continue to develop challenging, real-world datasets. By training on challenging, real-world data, detection algorithms will be more reliable in detecting firearms in real-world settings.

There are also firearm detection algorithm limitations relevant to public policy. For instance, criminal justice practitioners can only implement object detection algorithms in CCTV systems in public settings. Thus, the algorithms will not detect gun violence incidents on private property (e.g., homes, private businesses). Although this would limit firearm detectors' capabilities to an extent, the implementations in public settings will still be relevant to most gun violence events in the United States (Abt 2019).

Other limitations of firearm detection algorithms include the inability to decipher whether an individual with a firearm is a threat. Many localities, cities, and states across the United States permit the open carrying of a firearm. Although an individual who is legally open carrying will not break any laws or be a threat to society, the algorithms will still detect that firearm. Furthermore, as previously mentioned, the algorithms might produce false positives when objects resemble firearms. The threat of false positives is further compounded when the algorithms detect firearm replicas or toys (e.g., airsoft weapons, pellet guns). The inability to decipher who is a threat and the presence of false positives (especially of objects that closely resemble firearms) increases the likelihood that law enforcement is dispatched to non-dangerous scenes, potentially resulting in the mistaken use of force on an undeserving individual. Thus, practitioners must implement policies that help safeguard against misidentifying non-dangerous settings as dangerous (e.g., requiring a human operator to verify a firearm detection).

Additionally, some fear that the government's continued use of artificial intelligence (and general CCTV) will lead to abuses and violations of civil liberties, creating an authoritarian surveillance state (Feldstein 2019; Ratcliffe and Rosenthal 2021). Notwithstanding civil liberty concerns, firearm detection algorithms limit the risk of a surveillance state to an extent. Firearm detectors are only trained to detect guns. Nonetheless, although not relevant to firearm detection algorithms, it must be mentioned that other uses of object detection (e.g., facial recognition) might increase the likelihood of civil liberty intrusions and should be used carefully (and sparingly) (Zero Eyes n.d.).

Increased surveillance (via firearm detectors and CCTV) might also create biases in policing. Gun violence typically concentrates within minority and disadvantaged communities (Magee 2020). Thus, outside of implementations in schools or heavily traveled nodes, most implementations of firearm detectors (by police) might take place in those communities where gun violence is the deadliest. Not only will situating firearm detectors in minority and disadvantaged communities increase the surveillance of those residents, but it might also lead to more contacts with the police under false pretenses (e.g., due to an incorrect detection), which could lead to use of force disparities. As seen in the growing public health literature, police use of force can elicit varying adverse consequences for minority individuals and communities, including death (Simckes et al. 2021). Additionally, unjustified use of force might stem other issues, including lower trust in law enforcement, reduced police legitimacy, and lack of compliance with authority (Tyler 2004). Thus, oversight and accountability must be at the forefront when implementing artificial intelligence technologies that could increase violent interactions between the police and the public. Future work must balance the potential increase in public safety with possible intrusions into civil liberties to implement automated firearm detection interventions ethically and optimally.

Lastly, as previously mentioned, there has yet to be an empirical evaluation (to our knowledge) of firearm detection algorithms in a real-world setting. Despite situational crime prevention providing a theoretical mechanism to their potential efficacy, an experimental evaluation of firearm detection algorithms in real-world settings would provide greater confidence in their ability to prevent and mitigate gun violence. Although conducting a randomized controlled experiment to evaluate firearm detection algorithms presents practical challenges (given the statistical rarity and often unpredictable nature of gun violence) and raises ethical challenges, researchers and practitioners can take other steps to build an evidence base to determine the potential efficacy of firearm detection algorithms in security settings.

A simple evaluation might involve law enforcement and security organizations measuring precision and recall of firearm detection algorithms in real-world settings. For instance, embedding firearm detection algorithms in CCTV technologies in a designated area (away from the public) and exposing the CCTV viewshed to various negative (non-firearm) and positive (firearm) scenarios would provide a baseline understanding of how the technologies perform within CCTV systems in the real world. Future evaluations must also discern the performance disparities between humans and algorithms in firearm detection. Future evaluations should expose humans and algorithms to identical video data (ideally resembling dense CCTV footage) that contains various negative and positive scenarios. The comparison would yield data as to the differences between the algorithms' and humans' precision, recall, and time to detection when identifying firearms in video data. Other more resource-intensive evaluations might combine the previous two suggestions.

For example, an experiment could replicate real-world environments (e.g., crowds—through volunteers, research and practitioner personnel; poor visibility areas) and create various positive and negative scenarios to test performance differences (i.e., precision, recall, time to detection) between humans and firearm detectors embedded within CCTV systems. Although expensive, this type of experiment would provide the greatest understanding of how firearm detection algorithms can improve public safety and mitigate gun violence in real-world applications.

Notwithstanding associated limitations and concerns, object detection algorithms provide an exciting advancement to help solve the gun violence epidemic in the United States. Because object detection algorithms are only in their infancy (most advancements have come in the last 5 years), their implementation in CCTV technologies may not entirely supplant current practices. However, as mentioned above, detection algorithms' current capabilities make them ideal for assisting human operators in detecting objects of interest in live CCTV footage. The criminal justice system can better prevent or mitigate firearm-related events by embedding detection algorithms in CCTV systems and alerting human operators when a potential threat is present.

Conclusion

Gun violence takes the lives of tens of thousands of individuals in the United States each year. Although several fields have examined gun violence from their unique perspectives, less often do disciplines come together to put forth a holistic response. We set out to bridge the gap between the computer vision and criminology fields. In doing so, we present a theoretical framework that provides a causal mechanism to explain how and why firearm detection algorithms will work in real-world applications. Additionally, we demonstrate our approach to train an object detection algorithm and how scholars not previously exposed to computer vision can do the same. Our approach's results reaffirm that such algorithms can and will successfully detect firearms in real-world settings and assist human operators in monitoring dense CCTV live feeds. Although object detection algorithms are not without limitations, they are rapidly improving each year, and their current capabilities can more than improve the current uses of CCTV. Ultimately, gun violence is a multifaceted issue that affects public safety, and collaboration between varying disciplines is needed to create an effective solution.

Author contributions All authors have agreed for authorship, read and approved manuscript, and given consent for submission and subsequent publication of the manuscript. The paper properly credits the meaningful contributions of co-authors.

Funding This work is partially supported by the National Geospatial-Intelligence Agency.

Declarations

Conflict of interest The National Geospatial-Intelligence Agency (NGA) supported Tyler Houser while writing the manuscript. However, the NGA had no influence on this manuscript's content or outcome.

Thus, the authors declare no competing interests.

Ethical approval This manuscript is the authors' own original work. The manuscript does not include plagiarism, falsification of data, misuse of third-party material, or fabrication of results and fraudulent authorship. This study did not require IRB approval.

References

Abt, R. 2019. *Bleeding out: The devastating consequences of urban violence—And a bold new plan for peace in the streets*. New York: Basic Books.

Agnew, R. 1992. Foundation for a general strain theory of crime and delinquency. *Criminology* 30 (1): 47–87.

Ahmed, S., M.T. Bhatti, M.G. Khan, B. Lövström, and M. Shahid. 2022. Development and optimization of deep learning models for weapon detection in surveillance videos. *Applied Sciences* 12 (12): 5772. https://doi.org/10.3390/app12125772.

Akers, R.L. 1999. *Criminological theories: Introduction and evaluation*, 2nd ed. New York: Taylor & Francis.

Ashraf, A.H., M. Imran, A.M. Qahtani, A. Alsufyani, O. Almutiry, A. Mahmood, M. Attique, and M. Habib. 2022. Weapons detection for security and video surveillance using CNN and YOLO-V5s. *Computers, Materials and Continua* 70 (2): 2716–2775. https://doi.org/10.32604/cmc.2022.018785.

Berardini, D., L. Migliorelli, A. Galdelli, E. Frontoni, A. Mancini, and S. Moccia. 2023. A deep-learning framework running on edge devices for handgun and knife detection from indoor video-surveillance cameras. *Multimedia Tools and Applications*. https://doi.org/10.1007/s11042-023-16231-x.

Bhatti, M.T., M.G. Khan, M. Aslam, and M.J. Fiaz. 2021. Weapon detection in real-time CCTV videos using deep learning. *IEEE Access* 9: 34366–34382. https://doi.org/10.1109/ACCESS.2021.3059170.

Blaire, J.P., and K.W. Schweit. 2014. *A study of active shooter incidents, 2000–2013*. Texas State University and Federal Bureau of Investigation. https://www.fbi.gov/file-repository/active-shooter-study-2000-2013-1.pdf/view.

Braga, A.A., G.L. Pierce, J. McDevitt, B.J. Bond, and S. Cronin. 2008. The strategic prevention of gun violence among gang-involved offenders. *Justice Quarterly* 25 (1): 132–162. https://doi.org/10.1080/07418820801954613.

Buckchash, H., and B. Raman. 2017. A robust object detector: Application to detection of visual knives. In *2017 IEEE International Conference on Multimedia and Expo Workshops (ICMEW)*. https://ieeexplore.ieee.org/abstract/document/8026214/authors#authors.

Carson, A.E. 2021, April. Mortality in state and federal prisons, 2001–2018—Statistical tables (NCJ 255970). U.S. Department of Justice. https://bjs.ojp.gov/content/pub/pdf/msfp0118st.pdf.

CDC. 2022. *Firearm deaths grow, disparities widen.* CDC. https://www.cdc.gov/vitalsigns/firearm-deaths/index.html. Accessed 10 Jan 2023.

Clarke, R.V. 2017. Situational Crime Prevention. In Environmental Criminology and Crime Analysis, eds. Richard Wortley and Michael Townsley, 286–303. London and New York: Routledge.

Duran-Vega, M.A., M. Gonzalez-Mendoza, L. Chang, and C.D. Suarez-Ramirez. 2021. TYolov5: A temporal YOLOv5 detector based on quasi-recurrent neural networks for real-time handgun detection in video. https://arxiv.org/abs/2111.08867.

Federal Bureau of Investigation. 2022. *Active shooter incidents in the United States in 2021*. FBI. https://www.fbi.gov/file-repository/active-shooter-incidents-in-the-us-2021-052422.pdf/view.

Feldstein, S. 2019. The road to digital unfreedom: How artificial intelligence is reshaping repression. *Journal of Democracy* 30 (1): 40–52. https://doi.org/10.1353/jod.2019.0003v.

Freilich, J.D., S.M. Chermak, and B.R. Klein. 2020. Investigating the applicability of situational crime prevention to the public mass violence context. *Criminology and Public Policy* 19 (1): 271–293. https://doi.org/10.1111/1745-9133.12480.

Garza, J.I., and C.A. Vega. 2021. Gun detection in video frames with YOLOv3. https://www.researchgate.net/publication/356683717_Gun_Detection_in_Video_Frames_with_YOLOv3.

Goold, B.J. 2004. *CCTV and policing: Public area surveillance and police practices in Britain*. Oxford: Oxford University Press.

Gottfredson, M., and T. Hirschi. 2019. *Modern control theory and the limits of criminal justice*. Oxford: Oxford University Press.

Gu, Y., X. Liao, and X. Qin. 2022. YouTube-GDD: A challenging gun detection dataset with rich contextual information. https://arxiv.org/abs/2203.04129.

Guerette, R. 2009. The pull, push, and expansion of situational crime prevention evaluation: An appraisal of thirty-seven years of research. In *Evaluating crime reduction initiatives*, ed. J. Knutsson and N. Tilley, 29–58. Monsey: Criminal Justice Press.

Guerette, R.T., and K.J. Bowers. 2009. Assessing the extent of crime displacement and diffusion of benefits: A review of situational crime prevention evaluations. *Criminology* 47 (4): 1331–1368. https://doi.org/10.1111/j.1745-9125.2009.00177.x.

Hatten, D.N., and K.T. Wolff. 2020. Rushing gunshot victims to trauma care: The influence of first responders and the challenge of the geography. *Homicide Studies* 24 (4): 377–397. https://doi.org/10.1177/1088767920916900.

Idrees, H., M. Shah, and R. Surette. 2018. Enhancing camera surveillance using computer vision: A research note. *Policing: an International Journal of Police Strategies and Management* 41 (2): 292–307.

Katsiyannis, A., L.J. Rapa, D.K. Whitford, and S.N. Scott. 2023. An examination of US schools mass shootings, 2017–2022: Findings and implications. *Advances in Neurodevelopmental Disorders* 7: 66–76. https://doi.org/10.1107/s41252-022-00277-3.

Kaya, V., S. Tuncer, and A. Baran. 2021. Detection and classification of different weapon types using deep learning. *Applied Sciences* 11 (16): 7535. https://doi.org/10.3390/app11167535.

Kibria, S.B., and M.S. Hasan. 2017. An analysis of feature extraction and classification algorithms for dangerous object detection. In *2017 2nd International conference on electrical and electronic engineering (ICEEE)*, 2017. https://ieeexplore.ieee.org/abstract/document/8412846?casa_token=dCX_sYA0lHEAAAAA:6Pu78nHDAwqSIE0qK4Vk0HIsZnhCrQFMbFS9lLxBnSebXhSANVAPXIDtLwUgVYk3zOkdwcIJ.

Kolbe, L.J. 2020. School gun violence in the United States. *Journal of School Health* 90 (3): 245–253. https://doi.org/10.1111/josh.12866.

Koper, C.S., J. Egge, and C. Lum. 2015. Institutionalizing place-based approaches: Opening 'cases' on gun crime hot spots. *Policing* 9 (3): 242–254. https://doi.org/10.1093/police/pav023.

LeCune, Y., Y. Bengio, and G. Hinton. 2015. Deep learning. *Nature* 521: 436–444. https://doi.org/10.1038/nature14539.

Liang, K.J., J.B. Sigman, G.P. Spell, D. Strellis, W. Chang, F. Liu, T. Mehta, and L. Carin. 2019. Toward automatic threat recognition for airport X-ray baggage screening with deep convolutional object detection. https://doi.org/10.48550/arXiv.1912.06329.

Magee, L.A. 2020. Community-level social processes and firearm shooting events: A multilevel analysis. *Journal of Urban Health* 97 (2): 296–305. https://doi.org/10.1007/s11524-020-00424-y.

Makarios, M.D., and T.C. Pratt. 2012. The effectiveness of policies and programs that attempt to reduce firearm violence: A meta-analysis. *Crime and Delinquency* 58 (2): 222–244. https://doi.org/10.1177/0011128708321321.

McDonald, A. 2022, October 7. *Data quality considerations for machine learning models*. Towards Data Science. https://towardsdatascience.com/data-quality-considerations-for-machine-learning-models-dcbe9cab34cb#:~:text=Ensuring%20you%20have%20good%20data,are%20made%20on%20those%20results.

McKay, S., G.S. Hartnett, and B. Held. 2022. *Airline security through artificial intelligence: How the Transportation Security Administration can use machine learning to improve the electronic* *baggage screening program.* RAND Corporation. https://www.rand.org/content/dam/rand/pubs/perspectives/PEA700/PEA731-1/RAND_PEA731-1.pdf.

Narejo, S., B. Pandey, D.E. Vargas, C. Rodriguez, and M.R. Anjum. 2021. Weapon detection using YOLO v3 for smart surveillance system. *Mathematical Problems in Engineering*. https://doi.org/10.1155/2021/5808206.

National Threat Assessment Center. 2019. *Protecting America's Schools: A U.S. Secret Service analysis of targeted school violence*. United States Secret Service. https://www.secretservice.gov/sites/default/files/2020-04/Protecting_Americas_Schools.pdf.

Olmos, R., S. Tabik, and F. Herrera. 2017. Automatic handgun detection alarm in videos using deep learning. https://arxiv.org/abs/1702.05147.

Piza, E.L., and V.A. Sytsma. 2015. Exploring the defensive actions of drug sellers in open-air markets: A systematic social observation. *Journal of Research in Crime and Delinquency* 53 (1): 36–65. https://doi.org/10.1177/0022427815592451.

Piza, E.L., J.M. Caplan, L.W. Kennedy, and A.M. Gilchrist. 2015. The effects of merging proactive CCTV monitoring with directed police patrol: A randomized controlled trial. *Journal of Experimental Criminology* 11 (1): 43–69. https://doi.org/10.1007/s11292-014-9211-x.

Piza, E.L., B.C. Welsh, D.P. Farrington, and A.L. Thomas. 2019. CCTV surveillance for crime prevention: A 40-year systematic review with meta-analysis. *Criminology and Public Policy* 18 (1): 135–159. https://doi.org/10.1111/1745-9133.12419.

Qi, D., W. Tan, Z. Liu, Q. Yao, and J. Liu. 2021. A dataset and system for real-time gun detection in surveillance video using deep learning. https://arxiv.org/abs/2105.01058.

Ratcliffe, J., and J.M. Rosenthal. 2021. *Video surveillance of public places* (2nd ed.). ASU Center for Problem-Oriented Policing. https://popcenter.asu.edu/content/video-surveillance-public-places-2nd-edition. Accessed 5 Sep 2023

Reaves, B.A. 2015. *Local police departments, 2013: Equipment and technology*. U.S. Department of Justice.

Redmon, J., S. Divvala, R. Girshick, and A. Farhadi. 2015. You only look once: Unified, real-time object detection. https://arxiv.org/abs/1506.02640.

Rigano, C. 2019. *Using artificial intelligence to address criminal justice needs*. National Institute of Justice. https://www.ojp.gov/pdffiles1/nij/252038.pdf.

Rosales, M.P., C. Rodriguez, Y. Pomachagua, and C. Navarro. 2021, September 22–23, 2021. Firearm detection in images of video surveillance cameras with convolutional neural networks. In *13th International conference on computational intelligence and communication networks (CICN)*, Lima, Peru, 2021.

Rosenfeld, R., M.J. Deckard, and E. Blackburn. 2014. The effects of directed patrol and self-initiated enforcement on firearm violence: A randomized controlled study of hot spot policing. *Criminology (beverly Hills)* 52 (3): 428–449. https://doi.org/10.1111/1745-9125.12043.

Seiden, M. 2022, April 20. Almost 100 shanks made from crumbling pieces of Fulton County Jail seized from inmates. WSB-TV. https://www.wsbtv.com/news/local/fulton-county/almost-100-knives-made-pieces-deteriorating-fulton-county-jail-seized/VL7NLLQN2ZDT3DNX5KR3DOOMRI/.

Shah, D. 2022, March 7. Mean average precision (mAP) explained: Everything you need to know. V7 Labs. https://www.v7labs.com/blog/mean-average-precision#h1.

Sigman, J.B., G.P. Spell, K.J. Liang, and C. Lawrence. 2020. Background adaptive faster R-CNN for semi-supervised convolutional object detection of threats in X-ray images. *ADIX* 11404: 1140404. https://doi.org/10.1117/12.2558542.

Simckes, M., D. Willits, M. McFarland, C. McFarland, A. Rowhani-Rahbar, and A. Hajat. 2021. The adverse effects of policing on population health: A conceptual model. *Social Sciences* 281: 114103. https://doi.org/10.1016/j.socscimed.2021.114103.

Simon, T.R., et al. 2022. Notes from the field: Increases in firearm homicide and suicide rates—United States, 2020–2021. *Morbidity and Mortality Weekly Report* 71 (40): 1286–1287.

Skogan, W.G. 2019. The future of CCTV. *Criminology and Public Policy* 18 (1): 161–166. https://doi.org/10.1111/1745-9133.12422.

Snipes, J.B., T.J. Bernard, and A.L. Gerould. 2019. *Vold's theoretical criminology*. Oxford: Oxford University Press.

Tyler, T.R. 2004. Enhancing police legitimacy. *The Annals of the American Academy of Political and Social Science* 593 (1): 84–99. https://doi.org/10.1177/0002716203262627.

Wang, C., A. Bochkovskiy, and H.M. Liao. 2022. YOLOv7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors. https://arxiv.org/abs/2207.02696.

Zero Eyes. n.d.. Gun detection technology: 11 Common myths debunked. https://zeroeyes.com/gun-detection-technology-common-myths-debunked/. Accessed 5 Sep 2023.