Reducing Gun-Related Incidents on Construction Sites: An AI-Driven Approach for Automated Detection of People, Tools, and Firearms

Ammar Alzarrad, Ph.D.;[1] Shahid Ali;[2] Sudipta Chowdhury, Ph.D.;[3] and Husnu Narman, Ph.D.[4]

[1]Department of Civil Engineering, Marshall University, One John Marshall Drive, Huntington, WV 25755; e-mail: alzarrad@marshall.edu
[2]Department of Computer Sciences and Electrical Engineering, Marshall University, One John Marshall Drive, Huntington, WV 25755; e-mail: shahidali@marshall.edu
[3]Department of Mechanical and Industrial Engineering, Marshall University, One John Marshall Drive, Huntington, WV 25755; e-mail: chowdhurys@marshall.edu
[4]Department of Computer Sciences and Electrical Engineering, Marshall University, One John Marshall Drive, Huntington, WV 25755; e-mail: narman@marshall.edu

ABSTRACT

Traditional construction safety research primarily focuses on occupational hazards such as falls, equipment-related injuries, and environmental risks, while workplace violence, including firearm-related threats, is an underexplored area in construction safety research. High-profile cases, such as the tragic shooting on September 1, 2024, in Phoenix, Arizona, where a 48-year-old construction worker was fatally shot and his 19-year-old coworker sustained life-threatening injuries, underscore the importance of workplace violence mitigation in construction sites. To address firearm-related threats in construction, this study introduces RiskScanConstruction, a comprehensive dataset integrating the YouTube Gun Detection Dataset (YouTube-GDD) and the YOLO7 Power Tool Dataset. Designed specifically for construction safety applications, this dataset enables multi-class detection of people, firearms, and tools, distinguishing between similar-looking objects such as firearms and drills in dynamic, high-risk environments. The study evaluates the performance of various YOLO model configurations on RiskScanConstruction, with YOLOv11x emerging as the optimal model, achieving a precision of 92.0%, mAP@0.5 of 91.2%, and strong performance in detecting critical safety elements. The findings highlight the potential of advanced object detection technologies to proactively identify and mitigate construction site violence, enhancing both worker safety and site security. Furthermore, the RiskScanConstruction dataset and accompanying insights offer valuable contributions for broader safety-critical applications, including public security and industrial monitoring, setting the stage for future advancements in multi-class object detection systems.

INTRODUCTION

Construction sites are increasingly becoming settings for violent incidents, adding a new layer of risk to an already hazardous occupation (Fox 10 Phoenix, 2024). High-profile cases, such as the tragic shooting on September 1, 2024, in Phoenix, Arizona, have highlighted the vulnerability of workers in these environments. In that incident, a 48-year-old construction worker was fatally shot, and his 19-year-old coworker sustained life-threatening injuries following an altercation while they were repairing a fence. The attacker initially left the scene only to return with a firearm, ultimately leading to a fatal escalation (Fox 10 Phoenix, 2024). Such incidents are not isolated. Similar tragedies have occurred globally and underscore the urgent need for improved security and safety protocols on job sites (ABC News, 2023). From Auckland, New Zealand, where a construction worker fatally shot two colleagues in July 2023, to the violent shooting at Delaware County Linen near Philadelphia in May 2024, the frequency of such incidents is rising (The Philadelphia Inquirer, 2024).

Traditional safety measures in construction have primarily focused on mitigating physical risks, often overlooking the potential for violence among workers (Alzarrad et al., 2021). Recent statistics indicate a worrying rise in gun violence incidents, such as the shooting of a construction worker in Stockton, California, in October 2021 and an altercation at an Amazon facility in South Carolina in December 2022 that escalated to gunfire (Dillard, 2023). These incidents reveal that construction workers are vulnerable not only to physical hazards but also to violent confrontations that can lead to severe injuries or fatalities.

The need for a comprehensive approach to safety that includes strategies for preventing workplace violence is increasingly evident (Poulson et al., 2021). To effectively address these challenges, the integration of advanced technological solutions is essential. Object detection technologies, particularly those enhanced by machine learning and computer vision, can provide real-time situational awareness on construction sites (Hicks, 2023). However, existing object detection datasets, such as COCO and PASCAL VOC, are inadequate for the complex, dynamic environments typical of construction sites. They lack the necessary contextual richness and diversity of classes required for effective safety monitoring, particularly in detecting firearms and other critical safety elements (Richardson, 2023).

To bridge this gap, this study introduces two novel datasets tailored for construction safety applications: the YouTube Gun Detection Dataset (YouTube-GDD) and the RiskScanConstruction dataset. The YouTube-GDD is designed to enhance the detection of firearms in varied contexts, while RiskScanConstruction integrates this dataset with the YOLO7 Power Tool Dataset to enable multi-class detection of guns, tools, and people in dynamic construction environments (Jamieson & Romer, 2021). This comprehensive approach aims to evaluate various YOLO model configurations on the RiskScanConstruction dataset, providing insights into their effectiveness in detecting critical safety elements in real-time (Corburn et al., 2022).

This paper makes three key contributions: First, it introduces RiskScanConstruction, a robust dataset for safety-focused object detection, integrating guns, tools, and people in dynamic construction scenes. Second, it provides a comprehensive evaluation of several YOLO model variants on this dataset, including an analysis of class-specific performance. Third, it offers an insightful analysis of the trade-offs between model complexity, precision, and inference speed, delivering practical insights for deploying these models in safety-critical construction environments.

LITERATURE REVIEW

Object detection is a key technology for identifying and localizing objects within images or video streams, which is particularly important in safety-critical fields such as construction, manufacturing, and surveillance. In these environments, it enables real-time monitoring to prevent accidents and ensure regulatory compliance. For example, on construction sites, object detection aids in identifying heavy machinery, workers, and hazards, reducing the risk of collisions or falls (Huang et al., 2020; Zhang & Guo, 2020). In surveillance, object detection is critical for monitoring unauthorized access and protecting personnel in high-stakes environments (Gheisari & Esmaeili, 2019). Workplace violence, including gun-related threats, has become a rising concern in construction, where stressful conditions and tight deadlines may increase tensions (Roy & Islam, 2019). Object detection systems can identify unauthorized individuals and alert security personnel, bolstering site safety (Zhang & Guo, 2020). Implementing these technologies within safety management systems can enhance regulatory compliance and safety culture, reducing accidents and improving worker protection (Kim, 2018; Nouban & John, 2020; Mishra et al., 2022). The literature review section is divided into the following sections:

Existing Gun Detection Technologies

Gun detection technologies have progressed through advancements in machine learning and computer vision, particularly for surveillance. Current methods often employ deep learning, with convolutional neural networks (CNNs) being prominent for their capability to recognize firearms in real-time video feeds across diverse settings (Romero & Salamea, 2019; Srikar et al., 2023). These systems are generally trained on large datasets to identify guns under varied conditions, such as occlusions and different viewing angles (Rahman et al., 2022; Iqbal et al., 2021). Techniques such as HAAR-like feature-based detection in CCTV images show promise in controlled environments but may struggle in dynamic settings due to fluctuating lighting and background conditions (Rahman et al., 2022). Edge computing integration has enhanced real-time processing, facilitating faster responses to potential threats (Liu & HU, 2023). However, challenges remain, such as managing computational demands and achieving high accuracy in real-time scenarios (Liu & HU, 2023; Rahoo, 2023). While these technologies automate surveillance and reduce reliance on human operators, limitations include false positives and difficulty detecting firearms in unconventional positions (Srikar et al., 2023; Iqbal et al., 2021).

Gun Detection Datasets

Effective gun detection relies on high-quality datasets that reflect real-world complexities. The Gun-Detection-DB, a widely used dataset, contains images of firearms under varying conditions, including lighting and occlusions, providing foundational training data (Olmos et al., 2018; Srikar et al., 2023). However, datasets like Gun-Detection-DB primarily feature static images, lacking the dynamic contexts present in real-world scenarios, such as individuals in motion or crowded environments, which can limit model performance in complex situations (Moura et al., 2021). These datasets also often fall short in representing diverse firearm types and scenarios, thus impacting model robustness for broader practical applications (Iqbal et al., 2021). Inconsistencies in annotations further affect performance, particularly in cases requiring precise firearm localization (Moura et al., 2021). The need for comprehensive datasets that capture dynamic scenes, various firearm types, and contextual information is crucial for enhancing gun detection systems' effectiveness.

Dynamic Challenges in Gun Detection

Real-world gun detection is challenging due to variations in firearm orientation, partial occlusions, and the movement of individuals and objects in dynamic environments. Changes in firearm position affect visibility and recognition by detection algorithms, and occlusions often lead to false negatives (Srikar et al., 2023; Iqbal et al., 2021). Furthermore, movement in crowded spaces can blur images, complicating accurate processing (Gu et al., 2022). These challenges underscore the need for datasets that reflect dynamic scenarios, such as YouTube-Gun Detection Dataset (YouTube-GDD), which includes a range of contextual interactions to train adaptable algorithms (Gu et al., 2022). By incorporating diverse video scenarios, YouTube-GDD enables the study of firearm visibility under varying orientations and occlusions, promoting models that generalize well to real-world applications. The inclusion of contextual information, such as individual behaviors near a firearm, enhances detection accuracy when firearms are not fully visible, addressing a critical need for robust, real-world-ready gun detection technologies (Gu et al., 2022).

Motivation for Integrating YouTube-GDD and YOLO7 Power Tool Dataset

Integrating the YouTube-Gun Detection Dataset (YouTube-GDD) with the YOLO7 Power Tool Dataset offers a comprehensive safety solution for construction sites by combining each dataset's unique strengths. YouTube-GDD captures dynamic scenarios involving firearms, providing essential context for gun detection, while the YOLO7 Power Tool Dataset focuses on identifying power tools in construction settings. This integration enables a unified framework for multi-class detection—covering people, tools, and guns—enhancing situational awareness and safety monitoring on construction sites. A system capable of recognizing multiple classes in real-time allows for more efficient threat detection and quicker responses. For example, such a system can alert safety personnel if a nail gun is in use near someone carrying a firearm, heightening awareness and ensuring regulatory compliance. Moreover, training on a multi-class dataset improves model generalization to real-world scenarios by allowing it to learn interactions between objects, crucial for dynamic and cluttered construction environments where personnel and tools are in close proximity.

Unique Contributions of RiskScanConstruction Dataset

The RiskScanConstruction dataset advances construction safety object detection by addressing the limitations of single-class datasets. Traditional datasets often focus on isolated classes, limiting their applicability in multi-hazard environments like construction sites. RiskScanConstruction's inclusion of people, tools, and firearms enhances contextual relevance, allowing for multi-class detection in dynamic, cluttered settings. This capability improves situational awareness for safety managers, enabling prompt responses to potential hazards. Models trained on RiskScanConstruction benefit from learning object interactions, essential in settings where tools and personnel are in close proximity, such as a worker operating a nail gun near others. By balancing representation across classes, the dataset reduces overfitting, fostering robust models that generalize well to real-world conditions. Consequently, the RiskScanConstruction dataset offers a valuable, holistic resource for multi-class object detection systems, significantly improving safety outcomes in construction environments.

METHODOLOGY

Dataset Preparation

The RiskScanConstruction dataset was created by merging the YouTube Gun Detection Dataset (YouTube-GDD) and the YOLO7 Power Tool Dataset. Classes were unified into three categories: Class 0 (Person), Class 1 (Gun), and Class 2 (Tools). The dataset was divided into training (10,953 images), validation (1,213 images), and testing (1,127 images) sets, maintaining a balanced representation across classes. Annotation consistency was ensured through expert cross-validation, automated checks, and preprocessing steps, including standardizing image dimensions to 640x640 pixels. To enhance model generalization, data augmentation techniques such as random rotations, scaling, and flipping were applied. These augmentations simulated real-world conditions, improving robustness to variations in object positioning, orientation, and lighting.

Model Selection

The experiment evaluated eight YOLO model variants, each tailored to address specific detection challenges. YOLOv8n and YOLOv8x were designed as lightweight models, with the former optimized for edge devices and the latter for complex datasets. YOLOv9t and YOLOv9e incorporated enhanced feature pyramid networks to excel in detecting small and densely packed objects. YOLOv10n and YOLOv10x introduced transformer layers, enhancing contextual understanding and detection precision. Finally, YOLOv11n and YOLOv11x combined transformer networks with CNNs, offering high-resolution accuracy and adaptability, making them well-suited for safety-critical scenarios.

Training and Evaluation

Models were trained for 100 epochs with a batch size of 16, using early stopping to prevent overfitting. Evaluation metrics included precision, recall, mAP@0.5, inference speed, and F1 score, emphasizing accuracy and real-time performance critical for safety applications. Training logs were captured and analyzed to refine model performance. The training was conducted on a single GPU with comprehensive metric tracking and automated logging to ensure consistency and reliability in results. These efforts supported the development of a robust YOLO-based detection system tailored for construction safety monitoring. Figure 1 shows the research methodology flowchart.

Figure 1. Methodology Flowchart

RESULTS AND DISCUSSIONS

Overview of Results Across Models

The performance of each YOLO variant was assessed using precision, recall, mAP@0.5, inference speed, and F1 score. Among the tested models, YOLOv11x demonstrated the highest precision (0.920) and mAP@0.5 (0.912), indicating strong detection accuracy across all classes. YOLOv8x also showed high recall (0.878) and mAP@0.5 (0.914), suggesting it is well-suited for capturing a broad range of relevant instances. However, the inference speed for YOLOv11x and YOLOv8x was lower than lightweight models such as YOLOv8n and YOLOv9t, which provided faster processing but at the expense of slightly lower mAP and recall. Table 1 shows the performance across models.

Table1. Performance Metrics Summary Table

Model	Precision	Recall	mAP @ 0.5	Inference Speed (ms/frame)	F1 Score
YOLOv8n	0.885	0.838	0.889	0.667	0.861
YOLOv8x	0.904	0.878	0.914	5.883	0.891
YOLOv9t	0.891	0.831	0.900	0.915	0.861
YOLOv9e	0.898	0.859	0.905	7.247	0.878
YOLOv10n	0.883	0.822	0.888	0.699	0.852
YOLOv10x	0.906	0.850	0.907	5.514	0.877
YOLOv11n	0.898	0.834	0.894	0.647	0.864
YOLOv11x	0.920	0.850	0.912	5.747	0.884

Class-Specific Metrics and Visual Comparisons of Model Output

Class-specific performance analysis revealed that most models achieved higher precision and mAP for the “Person” class, while detection of the “Gun” class, which requires finer detail due to object variability, benefited from models with transformer layers, such as YOLOv10x and YOLOv11x. YOLOv9t and YOLOv11n, with faster inference speeds, showed a slightly lower mAP in detecting “Tools,” indicating trade-offs between speed and class-specific detection accuracy. Test images were used to visually compare the outputs of various models, highlighting YOLOv11x as the top performer. YOLOv11x consistently delivered robust detection in cluttered scenes, exhibiting fewer missed detections and high-confidence scores, particularly for small or partially obscured objects. In contrast, YOLOv8n, despite its faster processing speed, occasionally missed detections in complex images with overlapping tools and persons. Visual examples of YOLOv11x outputs demonstrate its superior accuracy in challenging scenarios. Figure 2 below shows samples of the models’ outputs.

Figure 2. Samples of Models Outputs

Precision and Recall

As shown in Figure 3, the bar graph comparison of precision and recall highlighted YOLOv11x as the top performer in precision (0.920), while YOLOv8x scored highest in recall (0.878). Models with high precision, like YOLOv11x, offered lower false positive rates, critical in applications where false alarms are costly. However, higher recall in YOLOv8x indicated better performance in comprehensive hazard detection, important for safety-critical settings where missing a detection could pose risks.

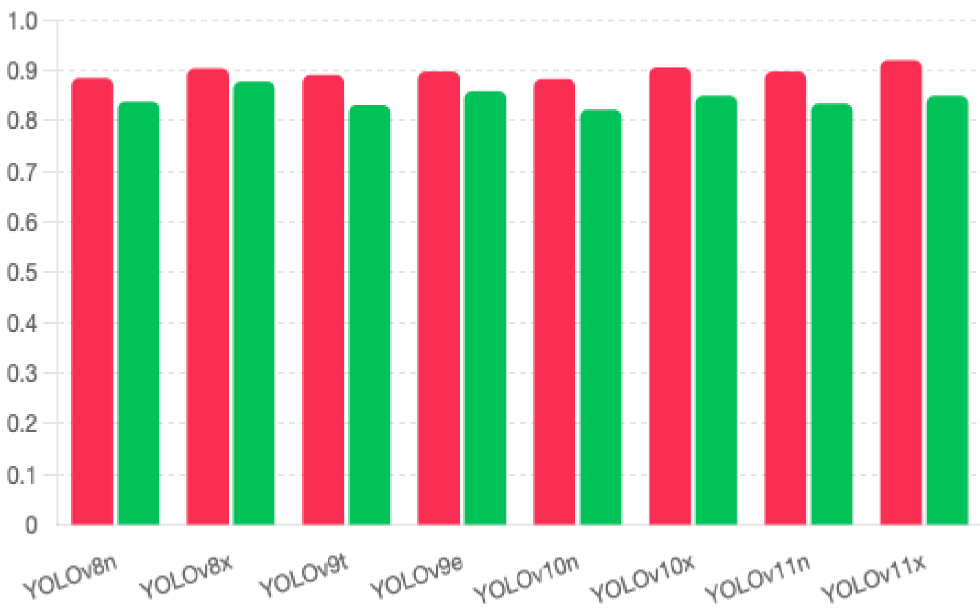

Figure 3. Bar Graph Comparison of Precision and Recall (Red for Precision and Green for Recall)

mAP@0.5 and Inference Speed

The line graph of mAP@0.5 across models showed YOLOv11x and YOLOv8x as the best-balanced models, each with mAP above 0.9, indicating strong detection performance across diverse classes. Models like YOLOv9e and YOLOv10x also performed well, though their slower inference speeds could impact real-time applications. A speed chart indicated that YOLOv8n (0.667 ms/frame) and YOLOv10n (0.699 ms/frame) achieved real-time speeds suitable for resource-limited environments. YOLOv11x and YOLOv8x, while slower, offered superior precision and mAP, making them more suited to scenarios where detection accuracy takes precedence over speed, as shown in Figure 4.

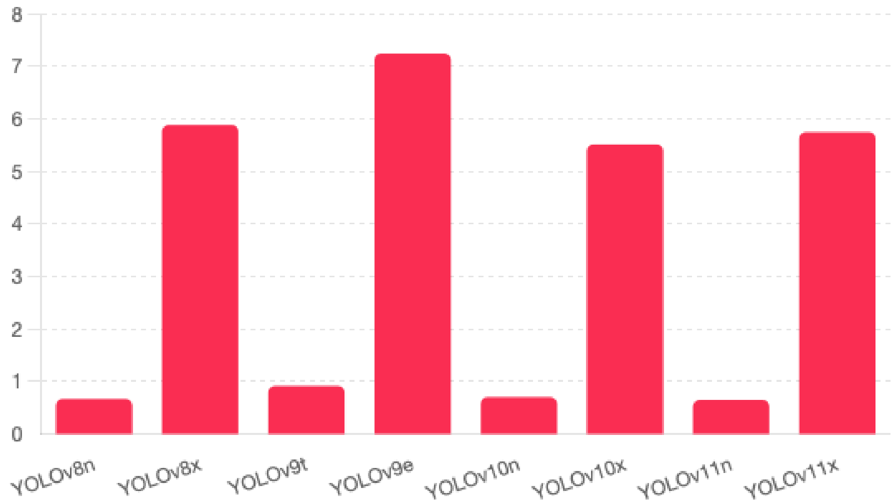

Figure 4. Inference Speed for Models

F1 Score and Optimal Model Selection

Class-specific F1 score analysis showed that YOLOv8x maintained consistent performance across classes, especially in detecting “Person” and “Gun” instances. YOLOv10x demonstrated balanced F1 scores for “Tools” and “Person” classes, indicating suitability for construction settings where accurate tool detection is necessary. Based on all metrics, YOLOv11x emerged as the optimal model, balancing high precision (0.920), mAP@0.5 (0.912), and competitive F1 score (0.884) with reasonable inference speed (5.747 ms/frame). While YOLOv8n and YOLOv9t offered faster speeds, YOLOv11x’s superior accuracy across classes and ability to capture intricate details make it well-suited for safety-critical applications that demand reliability over speed. Selecting YOLOv11x involves compromises, notably its slower inference speed compared to YOLOv8n. However, the improved detection quality justifies this trade-off in applications where accurate identification of potential hazards, such as guns and tools in construction, is prioritized over processing speed.

CONCLUSION

This study highlights the potential of advanced object detection technologies to enhance construction site safety by introducing RiskScanConstruction, a dataset specifically designed for detecting persons, firearms, and tools in dynamic, high-risk environments. The findings demonstrate that state-of-the-art YOLO models can effectively balance accuracy and computational efficiency, making real-time safety monitoring feasible. The results underscore the critical role of AI-driven detection systems in mitigating workplace hazards and preventing security threats. Despite these advancements, this research has certain limitations. The dataset, while comprehensive, does not yet account for a wider range of safety-related objects, such as helmets, safety barriers, and varied tool types, which are essential for broader safety assessments. Additionally, overlapping and occluded objects remain a challenge, affecting detection accuracy in cluttered environments. These limitations highlight the need for future research to expand object classes, integrate Vision Transformers (ViTs) and hybrid architectures, and refine real-time performance using techniques like model pruning and quantization. Addressing these challenges will enhance the scalability and deployment of AI-driven safety systems in construction and other high-risk industries.

REFERENCES

ABC News. (2023). 2 killed in New Zealand construction site shooting, suspect also dead. Retrieved from https://abcnews.go.com/International/2-killed-new-zealand-construction-site-shooting-suspect/story?id=101503562

Alzarrad, A., Emanuels, C., Imtiaz, M., and Akbar, H. (2021). "Automatic assessment of buildings' location fitness for solar panel installation using drones and neural networks." *CivilEng*, 2(1), 1052–1064. https://doi.org/10.3390/civileng2040056.

Corburn, J., Nidam, Y., & Fukutome-Lopez, A. (2022). The art and science of urban gun violence reduction: Evidence from the Advance Peace Program in Sacramento, California. *Urban Science*, *6*(1), 6. https://doi.org/10.3390/urbansci6010006

Dillard, D. (2023). Unlocking gun-violence solutions: The necessity and power of lived experience. *American Journal of Qualitative Research*, *7*(4), 182–202. https://doi.org/10.29333/ajqr/13729

Fox 10 Phoenix. (2024). Phoenix construction worker shot and killed, another injured while on the job. Retrieved from https://www.fox10phoenix.com/news/phoenix-construction-worker-shot-killed-another-injured-while-job

Gheisari, M., & Esmaeili, B. (2019). Applications and requirements of unmanned aerial systems (UASs) for construction safety. *Safety Science*, *118*, 230–240. https://doi.org/10.1016/j.ssci.2019.05.015

Gu, Y., Liao, X., & Qin, X. (2022). YouTube-GDD: A challenging gun detection dataset with rich contextual information. *arXiv preprint*. https://doi.org/10.48550/arxiv.2203.04129

Hicks, B. (2023). Who bought a gun during the COVID-19 pandemic in the United States? Associations with QAnon beliefs, right-wing political attitudes, intimate partner violence, antisocial behavior, suicidality, and mental health and substance use problems. *PLOS One*, *18*(8), e0290770. https://doi.org/10.1371/journal.pone.0290770

Holloway, K. (2023). Reviewing the literature on the impact of gun violence on early childhood development. *Current Psychiatry Reports*, *25*(7), 273–281. https://doi.org/10.1007/s11920-023-01428-6

Huang, Y., Hammad, A., & Zhu, Z. (2020). Providing proximity safety alerts to workers on construction sites using Bluetooth low energy RTLS. *Proceedings of the Creative Construction e-Conference*, *007*, 39–43. https://doi.org/10.3311/ccc2020-007

Iqbal, J., Munir, M. A., Mahmood, A., Ali, A. R., & Ali, M. (2021). Leveraging orientation for weakly supervised object detection with application to firearm localization. *Neurocomputing*, *440*, 310–320. https://doi.org/10.1016/j.neucom.2021.01.075

Jamieson, P., & Romer, D. (2021). The association between the rise of gun violence in popular U.S. primetime television dramas and homicides attributable to firearms, 2000–2018. *PLOS One*, *16*(3), e0247780. https://doi.org/10.1371/journal.pone.0247780

Kim, G. H. (2018). Measuring the effectiveness of safety incentives in construction sites in Korea. *Journal of Building Construction and Planning Research*, *6*(4), 267–277.

Kravitz-Wirtz, N., Bruns, A., Aubel, A., Zhang, X., & Buggs, S. (2022). Inequities in community exposure to deadly gun violence by race/ethnicity, poverty, and neighborhood disadvantage among youth in large U.S. cities. *Journal of Urban Health*, *99*(4), 610–625. https://doi.org/10.1007/s11524-022-00656-0

Liu, H., & Hu, Z. (2023). An edge computing-based handgun and knife detection method in IoT video surveillance systems. *International Journal of Advanced Computer Science and Applications*, *14*(11), 175. https://doi.org/10.14569/ijacsa.2023.0141117

Mishra, A. K., Adhikari, R., & Aithal, P. S. (2022). Linkage of safety site conditions with accidents. *International Journal of Health Sciences and Pharmacy*, *6*(1), 17–34. https://doi.org/10.47992/ijhsp.2581.6411.0077

Moura, N. S., Gondim, J. M., Claro, D. B., Souza, M., & Figueiredo, R. d. C. (2021). Detection of weapon possession and fire in public safety surveillance cameras. *Anais do Encontro Nacional de Inteligência Artificial e Computacional (ENIAC)*, 290–301. https://doi.org/10.5753/eniac.2021.18261

Nouban, F., & John, S. O. (2020). Developing a culture of health and safety on construction projects framework in Nigeria. *International Journal of Innovative Science and Research Technology*, *5*(7), 1019–1022. https://doi.org/10.38124/ijisrt20jul543

Olmos, R., Tabik, S., & Herrera, F. (2018). Automatic handgun detection alarm in videos using deep learning. *Neurocomputing*, *275*, 66–72. https://doi.org/10.1016/j.neucom.2017.05.012

Poulson, M., Neufeld, M., Dechert, T., Allee, L., & Kenzik, K. (2021). Historic redlining, structural racism, and firearm violence: A structural equation modeling approach. *The Lancet Regional Health - Americas*, *3*, 100052. https://doi.org/10.1016/j.lana.2021.100052

Rahman, S. U., Alam, F., & Ali, W. (2022). Gun detection in CCTV images using HAAR-like features. *Proceedings of Pakistan Academy of Sciences: A Physical and Computational Sciences*, *59*(4), 1–11. https://doi.org/10.53560/ppasa(59-4)749

Rahoo, A., Alvi, F. A., Rajput, U., & Halepoto, I. A. (2023). An efficient approach for firearms detection using machine learning. *VFAST Transactions on Software Engineering*, *11*(2), 94–99. https://doi.org/10.21015/vtse.v11i2.1467

Richardson, M. (2023). Understanding Latine perceptions of gun violence prevention strategies across New Mexico: A conceptual paper. *International Health Trends and Perspectives*, *3*(3), 352–364. https://doi.org/10.32920/ihtp.v3i3.1919

Romero, D., & Salamea, C. (2019). Convolutional models for the detection of firearms in surveillance videos. *Applied Sciences*, *9*(15), 2965. https://doi.org/10.3390/app9152965

Roy, C., & Islam, M. S. (2019). Hazards and safety issues at construction sites in Bangladesh. *MOJ Civil Engineering*, *5*(2), 52–56. https://doi.org/10.15406/mojce.2019.05.00149

Semenza, D. (2023). Reciprocal neighborhood dynamics in gun violence exposure, community health, and concentrated disadvantage in one hundred U.S. cities. *Journal of Urban Health*, *100*(6), 1128–1139. https://doi.org/10.1007/s11524-023-00796-x

Srikar, Y., Li, Z., Doshi, K. B., Mhasakar, P. M., Fan, H., Wei, J., Blasch, E., Zhang, B., & Ling, H. (2023). CCTV-Gun: Benchmarking handgun detection in CCTV images. *arXiv preprint*. https://doi.org/10.48550/arxiv.2303.10703

Stroud, A. (2020). Guns don't kill people…: Good guys and the legitimization of gun violence. *Humanities and Social Sciences Communications*, *7*(1). https://doi.org/10.1057/s41599-020-00673-x

The Philadelphia Inquirer. (2024). Shooting at Delaware County Linen leaves two dead. Retrieved from https://www.inquirer.com/news/live/chester-pa-shooting-deaths-victims-20240522.html

Zhang, Z., & Guo, H. (2020). A computer vision-based approach to classifying and storing image data for construction safety management. *Proceedings of the Creative Construction e-Conference*, *070*, 19–24. https://doi.org/10.3311/ccc2020-070