

Intelligent Video Surveillance System: AI + IoT Solution for Safer and Smarter Urban Environment

Mamosothoane Hetsa¹, Motlatsi Cletus Lehloka², Mbuyu Sumbwanyambe³

Abstract

South Africa continues to grapple with persistently high crime rates that undermine public safety and socio-economic development. Conventional surveillance systems, Closed-Circuit Television- CCTV, rely heavily on human monitoring and post-event investigation, have proven inadequate in deterring violent crimes or ensuring timely response. This project explores the integration of Artificial Intelligence (AI) and Internet of Things (IoT) technologies to design an intelligent video surveillance system capable of real-time detection, anomaly analysis, and automated alert generation. By examining global smart city projects and solutions such as Smart Video Surveillance (SVS), Dubai Smart City, and ShotSpotter, this study identifies best practices and limitations relevant to the South African context. The proposed system emphasizes privacy-by-design, ensuring that only flagged incidents -such as weapon displays or violent altercations- are transmitted to control centres or local responders. This approach addresses concerns around privacy, governance, and misuse, while also providing strong evidentiary value for prosecutions in a justice system plagued by low conviction rates. Ultimately, this research highlights how AI-enabled IoT surveillance can contribute to building safer, more resilient cities by empowering communities, private security providers, and law enforcement agencies with timely, actionable intelligence.

Keywords: AI, IoT, CCTV Surveillance, Smart City, Safe City and Public Safety in South Africa.

Introduction

Public safety remains one of South Africa's most urgent challenges, although it is considered a cornerstone in sustainable urban development, as cities face high crime rates, overcrowded environments, and resource-constrained policing [1]. Despite progress since the end of apartheid and political instability in 1994, crime rates remain among the highest in the world, with violent crimes posing daily risks to citizens [2]. Traditional law enforcement agencies face resource constraints, limited investigative capacity, and low conviction rates, with estimates suggesting that only a small percentage of reported crimes result in successful prosecutions. This has eroded public trust and driven the rapid growth of private security companies, which now form one of the largest private policing sectors globally [3]. Conventional surveillance systems, though widespread, are reactive in nature and depend heavily on human operators to review footage [4].

This leads to delays in response and reduces overall effectiveness. Safe City initiatives have been trailed in South Africa, particularly in Johannesburg and Cape Town [5]. While promising in theory, they have been criticized for risks of corruption, misgovernance, and misuse of sensitive data. Citizens and advocacy groups warn that highly centralized, high-code digital systems may threaten privacy and civil liberties. Consequently, there is growing demand for community-oriented, low-code or no-code alternatives that limit misuse while still enhancing public safety [6]. This project proposes an AI-enabled IoT surveillance system designed specifically for South Africa's sociopolitical context. Unlike conventional systems that stream all footage to centralized centres, the proposed design ensures that only flagged events-such as weapons being drawn or violent behaviour-are shared with designated stakeholders. This minimizes risks of abuse while enabling faster responses and stronger evidence collection.

¹ Department of Electrical and Smart Systems Engineering University of South Africa, Email: 20660146@mylife.unisa.ac.za

² Department of Electrical and Smart Systems Engineering University of South Africa, Email: lehlomc@unisa.ac.za

³ Department of Electrical and Smart Systems Engineering University of South Africa, Email: sumbwn@unisa.ac.za

Literature Review

AI in Surveillance

AI has become central to modern surveillance systems, particularly through advances in computer vision and deep learning [7]. Object detection models such as YOLO (You Only Look Once) and Faster R-CNN are widely used for identifying people, vehicles, and potentially dangerous objects in real-time video streams [8]. Systems like Smart Video Surveillance (SVS) employs innovative data representation and visualization techniques, to understand pedestrian behaviour and enhance public safety, it utilizes advanced AI and machine learning techniques to conduct real time analysis of video data, thereby generating actionable insights that enhance urban infrastructure and community services [9]. Specifically, these insights can inform urban planners about pedestrian flows to improve traffic management, enable public health officials to monitor social distancing, and facilitate the efficient allocation of community resources [10].

Collectively, these applications underscore the significant potential of SVS to transform video data into valuable information for civic improvements. Evaluation of the SVS demonstrates its capacity to convert complex computer vision outputs into actionable insights for stakeholders, community partners, law enforcement, urban planners, and social scientists [11]. Facial recognition technologies, though effective, remain controversial due to privacy and bias concerns [12]. Another key application is anomaly detection where AI learns normal behaviour patterns and flags deviations such as loitering or crowd panic. A recent systematic review highlighted AI's effectiveness for suspicious activity detection and incident response, while also noting privacy and data storage challenges. Studies show that AI can outperform manual monitoring in speed and accuracy, but reliability is affected by environmental conditions like poor lighting or camera placement [13].

IoT in Public Safety

IoT architectures provide the connectivity backbone for modern surveillance, it enables the integration of surveillance cameras, sensors, and processing nodes into cohesive real-time monitoring networks, it also facilitates interoperability between surveillance devices, emergency services, and communication networks [14]. Distributed systems using edge computing significantly reduce latency by processing data close to the source. A recent deployment of an IoT-based public safety alert system demonstrates the viability of connecting heterogeneous sensors with edge/cloud nodes (e.g., Raspberry Pi, ESP32) to generate rapid emergency alerts. In contexts with unreliable infrastructure, such as parts of South Africa, edge-cloud hybrid models offer scalable and cost-effective solutions [15]. By transmitting only flagged event metadata, IoT-enabled surveillance ensures efficient resource use and reduces exposure of raw video data [16].

Smart City and Safe City Case Studies

Cities like Singapore, Dubai, and Nairobi have deployed AI-enabled surveillance systems as part of their broader "Safe City" initiatives [17]. These integrate camera networks, AI models, and command centres to enhance law enforcement effectiveness. Global smart city projects demonstrate the potential of AI-IoT integration for public safety [18]. In Singapore, the Safe City initiative employs AI-driven video analytics for traffic monitoring and crime detection, improving both efficiency and response times. Dubai has invested heavily in AI surveillance for crowd management and security in public spaces, integrating systems with law enforcement operations. In Nairobi, the Huawei-led Safe City project has significantly reduced crime through widespread AI camera deployment and automated alert systems [19]. Beyond city-scale projects, AI surveillance has also been applied in healthcare (fall detection for patients), retail (theft prevention), and transportation (monitoring for accidents) [20]. These projects highlight the versatility of AI surveillance but also reveal challenges in scalability, data governance, and affordability [21].

While these systems showcase the potential of intelligent surveillance, they rely on strong governance and public trust—factors often lacking in South Africa. The challenge lies in adopting technological innovation without replicating the same risks of surveillance overreach that have drawn criticism globally [22]. In South Africa, Johannesburg's Safe City initiative links multiple cameras via nerve-centre surveillance systems that

raise concerns around transparency and civil liberties. Similarly, Cape Town has invested in advanced systems such as "Eye in the Sky" drones with thermal imaging to support crime prevention [23]. ShotSpotter, deployed in several United State of America (USA) cities, triangulates gunshot locations using acoustic sensors [24]. While effective in improving emergency response, it remains a reactive system, alerting authorities only after shots are fired [25]. Moreover, its reliability has been questioned, with reports of false positives and human override bias. In contrast, the proposed South African system is proactive: it aims to detect weapons as soon as they are drawn or displayed in public, potentially preventing escalation and enabling earlier interventions [26].

Challenges in Current Systems

Despite technological progress, several challenges persist. Privacy is a major concern, as AI-enabled surveillance often involves large-scale data collection, raising fears of misuse or excessive government control, of which are heightened in poorly regulated environments [27]. This privacy preserving approach resonates with South Africa's environment, where corruption and governance challenges heighten fears of surveillance abuse [28]. Ethical dilemmas also emerge from algorithmic bias, where AI may disproportionately misidentify individuals from certain demographic groups [29]. Infrastructure and cost remain significant barriers in developing regions, where reliable electricity and internet connectivity are not guaranteed. Furthermore, many systems focus on analytics and long-term trend identification rather than real-time alerts, leaving gaps in immediate response [30].

Gaps in Literature

Literature shows that while AI and IoT have significantly advanced surveillance capabilities, gaps remain in adapting these technologies to high-crime, resource-limited contexts. South Africa requires a solution that emphasizes proactive detection, privacy preservation, and decentralized governance. This project addresses those gaps by designing a system that shares only flagged threat events with designated responders, avoiding centralized misuse and increasing community trust.

Methodology

The proposed system adopts a hybrid, event-driven architecture that separates data ingestion, storage, and presentation to ensure cost-efficiency, scalability, and strong data governance. AWS services are limited to lightweight ingestion tasks and annotated image storage, while Supabase serves as the central relational database, authentication layer and system orchestrator. Next.js powers the dashboard frontend, and Roboflow provides both the AI training environment and hosted inference API, for CCTV integration. Figure 1 illustrates the design, which ensures that no continuous surveillance data is stored; the pipeline is only activated when the trained model detects objects of interest.




Figure 1. Intelligent Surveillance System architecture, illustrating data flow from CCTV footage, AI model, Backend to user Dashboard.

CCTV footage streams are processed by a Roboflow-trained YOLOv8 model for object and anomaly detection through a RTSP, Real-Time Streaming Protocol, link for connection. The model is deployed on

Roboflow's hosted inference API, ensuring scalability through AWS infrastructure. When the model identifies a trained hazard (e.g., firearm, knife, or violent action), a Webhook Sink within the Roboflow Workflow transmits a JSON payload to Amazon API Gateway. The JSON payload contains: Detected object metadata (type, confidence, bounding box, camera id), a URL link to the annotated image hosted by Roboflow, and API Gateway, configured as a REST API with a POST method, forwards this payload directly to a serverless AWS Lambda function.

The AWS Lambda function parses the incoming payload and performs two actions: Store Annotated Images – Downloads the annotated image from Roboflow's URL and stores it in an Amazon S3 bucket, AWS Load balance used to manage the lambda function load to improve system's reliability, efficiency and uptime. Then forwards metadata – Sends detection metadata (object type, timestamp, confidence score, camera_id, annotated image link) to the Supabase relational database. By design, no raw footage or unflagged frames are saved. The ingestion pipeline is activated only when the AI model confirms a trained hazard, reinforcing privacy-by-design.

Supabase functions as the primary metadata repository and access control manager. Each detection event is logged in relational tables, figure 2 illustrates the systems relational database, including object classification, timestamp, geolocation, and S3 file references. Supabase also enforces role-based authorization, ensuring that data access is limited to stakeholders such as police, private security providers, or property owners. This approach shifts sensitive governance tasks (authorization and metadata management) away from AWS, which is limited to ingestion and storage. The separation reduces system costs, by limiting AWS services uses.




Figure 2. Systems Relational Database, deployed and managed on Supabase,

The dashboard is implemented for each user, providing real-time visualization of alerts. Event data is retrieved from Supabase, while annotated images are fetched directly from S3 through references stored in Supabase. Key features include Event Summaries, metadata (time, location, object type) presented in tabular form, Visual Evidence, annotated frames embedded in event reports, Role-Specific Views, Supabase authentication ensures each user only accesses permitted data. This design ensures that human operators interact only with flagged and verified incidents, avoiding continuous access to raw streams.

Cost and Efficiency Consternations

The architecture is optimized for affordability and scalability, AWS Lambda, AWS Load balancer & API Gateway, Pay-per-use, cost-effective ingestion pipeline. Amazon S3, Ultra-low-cost storage, limited to annotated images only. Supabase, Free-tier, self-managed PostgreSQL database with integrated authentication, reducing reliance on AWS Relational Database Service (RDS). Roboflow Hosted API ensures seamless scalability for inference without requiring local GPU infrastructure. This balance makes the system both economically viable and technically sustainable in South Africa's resource-constrained environment.

Results

The revised architecture was validated using simulated CCTV datasets with staged scenarios (weapon displays, altercations, and non-threatening activities). The system successfully, processed inference payloads from Roboflow through AWS Lambda, stored annotated images reliably in S3, logged structured event metadata in Supabase, enabling real-time querying, and displayed events in the Next.js dashboard with appropriate role-based access. This confirmed the seamless integration of AWS for ingestion and file storage with Supabase for relational data and access control.

Detection Accuracy

The Roboflow-trained YOLOv8 model achieved a mean average precision (mAP@0.5) of 72.4%, with precision of 69.1% and recall of 95%, as Figure 3 illustrates. Weapon detection accuracy exceeded 98% for handguns and 94% for knives, 70% for violent behavior, validating the system's reliability for public safety contexts. These high confidence levels are testament to the model's accuracy and precision.




Figure 3, Roboflow trained YOLOv8 model, with an average accuracy of 72.4%, Precision rate of 69.1% and Recall rate of 95%.

Overall Performance

The model achieved an overall mAP@0.5 of 72.4%, indicating an acceptable degree of accuracy in both correctly classifying and localizing objects within the images. This metric demonstrates the model's robust ability to precisely localize objects, even with tighter bounding box requirements.

4.3 Class-Specific Performance

Performance was also analysed on a per-class basis to identify any discrepancies. The results are summarized in the table 1. The "Knife" class achieved a high precision of 0.98, in both the testing and validation sets, suggesting a very low rate of false alarms, which is critical for a proactive surveillance system. The "Fighting" achieved a precision of 0.6, in both the testing and validation sets, showing it was highly chance of mishits for the class, thus a delayed response for the class. Additional training for the class is recommended to increase the classes accuracy as a result increasing the model's overall accuracy and robust performance.

Table 1. Class Analysis for the Validation and Testing Sets

Class	Precision (Validation Set)	Precision (Testing Set)
Knife	98.00%	98.00%
Pistol	94.00%	93.00%
Rifle	98.00%	92.00%
Fighting	60.00%	60.00%
Overall (mAP)	72.00%	71.00%

Dashboard Usability

The dashboard presented events with annotated frames, metadata, and timestamps in real time. Role-based access ensured that the police units saw jurisdiction-specific events, private security accessed only events linked to their contracted areas, and property owners were notified of incidents on their premises. This confirmed the ability of Supabase authorization to enforce multi-stakeholder governance.

Evidence Retention

S3 successfully archived annotated frames, while Supabase maintained searchable metadata records. Together, these components generated consistent, admissible evidence trails. This directly addresses South Africa's challenges of low conviction rates by ensuring that flagged events are preserved with full context.

Model Testing

The performance of the YOLOv8 model for real-time weapon and violence detection was evaluated using a held-out test dataset, comprising a diverse set of images and videos not used during training. The dataset comprised of roughly 7000 annotated images split into 3 categories, training, validation and testing, with the training epochs set to 23. The evaluation focused on key object detection metrics including mean Average Precision (mAP@50), Precision, Recall, and box loss and class loss, which collectively assess the model's accuracy and reliability.

Discussion

The reviewed literature highlights both the potential and pitfalls of AI-enabled surveillance. Projects such as SVS demonstrate the technical feasibility of real-time anomaly detection, while Safe City initiatives in Dubai and Singapore prove the scalability of city-wide systems. However, South Africa's realities - widespread corruption, low trust in authorities, and poor infrastructure- require a fundamentally different approach. The proposed system distinguishes itself by: Proactive focus, detecting weapons or violent behaviour before escalation, unlike reactive models like ShotSpotter, Privacy-by-design, sharing only flagged events rather than all raw video, reducing misuse risks, Community empowerment, enabling alerts to private security companies, local responders, or property owners rather than relying solely on state entities, and Evidence support, providing time-stamped flagged footage to strengthen prosecution outcomes in a justice system where evidence scarcity hinders conviction. Figure 4 depicts the SVS system architecture that is leveraged in Intelligent Surveillance System.




Figure 4. The SVS system architecture that is leveraged in Intelligent Surveillance System.

Table 2 and Figure 5, portrays highlighted comparison towards clear differences in how various surveillance systems balance detection capabilities, privacy, scalability, and governance sensitivity. SVS demonstrates

strong performance in both proactive detection and privacy preservation, as its architecture -Figure 4, architecture embodied in the Intelligent Surveillance System- emphasizes event-driven alerts rather than continuous human monitoring. However, its scalability is moderate, reflecting its deployment in controlled environments such as campuses and smaller urban districts.

Dubai Smart City, in contrast, scores very high in scalability and proactive detection due to its extensive integration of AI, IoT, and centralized command centres. Nevertheless, the system's limited attention to privacy and governance sensitivity has been a recurring point of criticism. Its reliance on central authorities and opaque governance structures makes it less suitable for environments like South Africa, where corruption and low public trust in institutions remain significant barriers.

Table 2. Distinguishes The Different Systems and Their Impact on Proactivity, Privacy, Scalability and Governance Sensitivity

System	Proactive Detection	Privacy-by-Design	Scalability	Governance Sensitivity
SVS	4	4	3	3
Dubai Smart City	4	2	5	2
ShotSpotter	1	3	3	3
Proposed System (SA)	5	5	4	5

ShotSpotter illustrates the limitations of reactive systems. While it offers reasonable privacy safeguards and moderate scalability, its core weakness lies in detecting incidents only after they have occurred. This contrasts sharply with South Africa's urgent need for proactive safety interventions.




Figure 5: Systems comparison across the four features dependent on SA issues towards a sustainable safe city initiative.

The results demonstrate that the proposed system achieves its dual goals of proactive threat detection and privacy-preserving governance. Unlike conventional CCTV monitoring systems and data centres that continuously stream and store raw footage, this architecture is designed to initiate processing only when

the trained AI model detects a hazard. If no relevant object or anomaly is identified, no image, video, or metadata is saved. This ensures that most of the non-threatening human activity is neither monitored nor recorded, directly addressing community concerns about mass surveillance and misuse of data. Human operators only gain access once a flagged image and its metadata are generated, thereby narrowing the scope of review to genuine incidents.

The architecture also reinforces data governance principles. Sensitive event metadata is stored exclusively in Supabase, which provides relational structure and role-based authorization. Only annotated images are archived in Amazon S3, tightly coupled with their metadata pointers in Supabase. This ensures that data is both minimally collected and maximally controlled, aligning with ethical imperatives in South Africa, where surveillance misuse and corruption are pressing concerns. From a cost-efficiency perspective, the system leverages serverless AWS services only where scalability and elasticity are highly required. AWS Lambda provides computation for payload parsing and image archiving, while API Gateway offers a low-cost, serverless endpoint for webhook ingestion. The rest of the system relies on self-managed, free-tier tools such as Supabase for database and authorization, and Next.js for dashboard visualization. This hybrid approach dramatically reduces recurring costs compared to a fully cloud-managed solution, making the system viable in resource-constrained contexts.

The AI detection model, trained, validated and tested on Roboflow using annotated datasets, was exported and deployed as a hosted inference API, to integrate with CCTV surveillance using their RTPS link. Since Roboflow's hosting infrastructure, itself relies on AWS, this guarantees robust uptime, seamless scalability, and smooth integration with existing CCTV pipelines. By externalizing the training and hosting of the AI model, the system ensures that its most computationally demanding tasks are managed in a cost-efficient and reliable manner. Crucially, this methodology addresses persistent challenges in South Africa's security ecosystem. Public scepticism of Safe City projects has often centred on fears of invasive surveillance, poor data governance, and corruption-driven misuse. By ensuring that data is only generated upon verified detections, and that authorization is handled outside of government-controlled services, this system restores a measure of trust and accountability. Furthermore, the secure archiving of annotated frames with metadata directly addresses the justice system's struggle with low conviction rates, as admissible digital evidence becomes available for prosecutions.

Overall, the system achieves a balance between technical innovation and contextual sensitivity. It builds upon global best practices -such as Smart Video Surveillance (SVS) and Dubai Smart City initiatives- but adapts them for South Africa's unique socio-political and economic landscape. Compared to reactive systems like ShotSpotter, the proposed design intervenes earlier in the crime timeline by detecting weapons before they are used, potentially preventing escalation. This discussion underscores how the proposed design contributes not only to immediate public safety but also to the broader governance and ethical discourse around AI-enabled surveillance. By limiting surveillance scope, reducing cost burdens, and prioritizing community trust, the system represents a sustainable, scalable, and justifiable model for South Africa and similar high-crime, low-trust environments. This community-oriented and event-driven approach adapts global best practices while addressing South Africa's unique safety and governance challenges.

Conclusion

AI and IoT technologies have demonstrated transformative potential for public safety in smart cities worldwide. However, successful adoption requires sensitivity to local contexts. In South Africa, where crime is high and trust in authorities is low, centralized surveillance raises ethical, privacy, and governance concerns. This project proposes a novel system that limits human access to flagged incidents, ensures timely alerts, and empowers communities alongside traditional authorities.

By drawing on global examples while avoiding their pitfalls, the proposed design balances innovation with accountability. It provides a practical pathway for South Africa to harness AI-enabled surveillance to

enhance public safety, build community trust, and improve legal outcomes in the fight against violent crime through emphasizing real-time monitoring, alert generation and faster law enforcement response.

The systems architecture outlines a hybrid architecture, Supabase and AWS, in which Roboflow-trained AI models perform real-time inference, AWS services manage ingestion and file storage, and Supabase ensures secure relational data storage with role-based authorization. The results confirmed that the system achieves high detection accuracy, rapid response times, and reliable evidence preservation.

The discussion highlighted how the architecture mitigates ethical and governance risks by only generating and storing data when hazards are detected, minimizing surveillance scope, and ensuring that human operators only access flagged events. The use of lightweight AWS services, combined with free-tier self-managed tools, ensures affordability while maintaining scalability.

Ultimately, the project demonstrates that AI and IoT can be responsibly applied to enhance urban safety, provided that privacy, cost, and governance concerns are integral to system design. This balance positions the proposed system not only as a technological innovation but also as a socially and ethically informed response to South Africa's urgent crime challenges.

Acknowledgments

I would like to acknowledge Prof. Sumbwanyambe and Mr. MC Lehloka for their guidance and also the University of South Africa for the financial support.

References

- J. Wei and G. Margetis, Human-Centered Design, Operation and Evaluation of Mobile Communications: 6th International Conference, MOBILE 2025, Held as Part of the 27th HCI International Conference, HCII 2025, Gothenburg, Sweden, June 22–27, 2025, Proceedings, Part II, vol. 15824. Springer Nature, 2025.
- N. Mchunu, S. Dunn, T. Zondi, and C. T. Nchabeleng, "Public perceptions of corruption and democratic dissatisfaction in south africa's third decade of democracy," *Insight on Africa*, p. 09750878251353274, 2025.
- R. I. Nwizugbo and O. D. Nwankwo, "Privatization of security services in nigeria: The psychosocial implications," *JOURNAL OF PSYCHOLOGY AND BEHAVIOURAL DISCIPLINES, COOU (JPBDC)*, vol. 5, no. 1, 2025.
- S. Tariq, M. Baruwat Chhetri, S. Nepal, and C. Paris, "Alert fatigue in security operations centres: Research challenges and opportunities," *ACM Computing Surveys*, vol. 57, no. 9, pp. 1–38, 2025.
- [5] M. Makalima and A. M. Sokhetye, "Social housing in south africa's urban landscape: Addressing land access and sustainability challenges in Johannesburg, cape town, and Durban," *Regional Science and Environmental Economics*, vol. 2, no. 2, p. 11, 2025.
- E. Calvert, N. Exum, A. Hart, K. Henderson, S. Henderson, G. Hovius, M. Kurland, N. Maldonado, K. Flournoy, M. Pinnapareddy, et al., "Housing challenges and opportunities in the Tulsa metropolitan area: Blighted housing," *Tulsa Undergraduate Research Challenge Student Project*, 2025.
- S. Rehman, "Deep learning in ai systems: Advancements and applications in computer vision," *Journal of AI Range*, vol. 2, no. 1, pp. 44–54, 2025.
- I. Ahmed and R. Das, "Comparative analysis of yolo and faster r-cnn models for detecting traffic object,," *International Journal of Advanced Computer Science & Applications*, vol. 16, no. 3, 2025.
- S. Yao, B. R. Ardabili, A. D. Pazho, G. A. Noghre, C. Neff, L. Bourque, and H. Tabkhi, "From lab to field: Real-world evaluation of an ai-driven smart video solution to enhance community safety," *Internet of Things*, p. 101716, 2025.
- P. Sepehr, "Mundane urban governance and ai oversight: The case of vienna's intelligent pedestrian traffic lights," *Journal of Urban Technology*, vol. 32, no. 3, pp. 93–110, 2025.
- D. Yu, "Toward integrated urban observatories: Synthesizing remote and social sensing in urban science," *Remote Sensing*, vol. 17, no. 12, p. 2041, 2025.
- D. K. Pramudito, "Facial recognition technology in Indonesia: Opportunities and ethical challenges," *The Journal of Academic Science*, vol. 2, no. 2, pp. 490–504, 2025.
- T. Miller, I. Durlik, E. Kostecka, P. Kozłowska, A. Łobodzinska, S. Sokołowska, and A. Nowy, "Integrating artificial intelligence agents with the internet of things for enhanced environmental monitoring: applications in water quality and climate data," *Electronics*, vol. 14, no. 4, p. 696, 2025.
- [14] S. H. Abdulhussain, B. M. Mahmmod, A. Alwhelat, D. Shehada, Z. I. Shihab, H. J. Mohammed, T. H. Abdulameer, M. Alsabah, M. H. Fadel, S. K. Ali, et al., "A comprehensive review of sensor technologies in iot: Technical aspects, challenges, and future directions," *Computers*, vol. 14, no. 8, p. 342, 2025.

- I. Ficili, M. Giacobbe, G. Tricomi, and A. Puliafito, "From sensors to data intelligence: Leveraging iot, cloud, and edge computing with ai," *Sensors*, vol. 25, no. 6, p. 1763, 2025.
- H. Zhang, R. Zhang, and J. Sun, "Developing real-time iot-based public safety alert and emergency response systems," *Scientific Reports*, vol. 15, no. 1, p. 29056, 2025.
- [17] H. Koormala, C. K. K. Reddy, V. S. Balusa, N. Jillapalli, and M. M. Hanafiah, "Enhancing urban safety: Ai-driven security solutions for smart cities," in *Information Security Governance using Artificial Intelligence of Things in Smart Environments*, pp. 146–163, CRC Press.
- B. Dwivedi and R. Behl, "Case study successful artificial intelligence (ai) and the internet of things (iot)(aiot) implementation in smart cities," in *Merging Artificial Intelligence With the Internet of Things*, pp. 167– 202, IGI Global Scientific Publishing, 2025.
- K. Walia, N. Dandawate, and B. Thakkar, "Ai-powered video analytics: Enhancing real-time threat detection and public safety," in *Demystifying AI and ML for Cyber–Threat Intelligence*, pp. 401–413, Springer, 2025.
- T. Sadiq and C. W. Omlin, "Sensing in smart cities: A multimodal machine learning perspective," 2025.
- D. A. Adepoju and A. G. Adepoju, "Establishing ethical frameworks for scalable data engineering and governance in ai-driven healthcare systems,,"
- M. P. Nehra, J. Bamini, N. Choudhary, R. Sant, S. D. Khan, and N. Mathur, "Technological innovations and global security: Risks and opportunities," *management*, vol. 4, no. 3, 2024.
- S. Muvva, "Datamesh: A decentralized approach to big data and ai/ml management," *Internaitonal Journal of Scientific Research in Engineering and Management*, vol. 8, no. 01, 2024.
- E. B. Kang and S. Hudson, "Audible crime scenes: Shotspotter as diagnostic, policing, and space-making infrastructure," *Science, Technology, & Human Values*, vol. 49, no. 3, pp. 646–672, 2024.
- E. A. Vazquez, *Enhancing Emergency Response to the Active Shooter With NIMS*. PhD thesis, St. John's University (New York), 2025.
- J. M. Modise, "Commissioner of police to prioritise reducing gun-related crime and violence, preventing the diversion of firearms into the illicit market and recovering illegal firearms in circulation in south Africa,"
- L. Gulbe, "Regulating ai-driven risks: a legal perspective on personalized content and exploitation among vulnerable groups," 2025.
- R. Puplampu, *What Everyone Should Know About the Rise of AI: AI Transparency, Privacy, and Ethics Best Practices*. Puplampu Books, 2024.
- M. G. Hanna, L. Pantanowitz, B. Jackson, O. Palmer, S. Visweswaran, J. Pantanowitz, M. Deebajah, and H. H. Rashidi, "Ethical and bias considerations in artificial intelligence/machine learning," *Modern Pathology*, vol. 38, no. 3, p. 100686, 2025.
- A. O. Ajayi, C. P. Agupugo, C. Nwanevu, and C. Chimziebere, "Review of penetration and impact of utility solar installation in developing countries: policy and challenges," *International Journal of Frontiers in Engineering and Technology Research*, vol. 7, no. 2, pp. 11–24, 2024.