International Journal of Computer Techniques – IJCT
Volume 13 Issue 4, August 2026
Open Access and Peer Reviewed Journal ISSN 2394-2231
https://www.ijctjournal.org/

RESEARCH ARTICLE
OPEN ACCESS
RESEARCH ARTICLE
OPEN ACCESS

AI-Driven Surveillance Auditing to Police Ethically

Surya Prakash Guna Sekhar*, Karthik Sai Pothireddy**
*(Department of Computer Science, University of Illinois, Springfield, United States
Email: suryaprakashgunasekhar@gmail.com)
**(Department of Operations & Management Information Systems, Northern Illinois University, DeKalb, United States
Email: karthiksaipothireddy@gmail.com)

Abstract: **The growth of Body-Worn Cameras (BWCs) and CCTV systems in law enforcement has generated unprecedented volumes of surveillance data, presenting new opportunities for enhancing transparency, accountability, and public trust. However, manual review of this footage remains slow, inconsistent, and often significantly delayed, creating substantial barriers to timely oversight. This study addresses this gap by designing an AI-driven surveillance auditing system capable of detecting unethical practices in policing — including acts of aggression, racial profiling, and procedural violations — to strengthen transparency, accountability, and public trust. Multimodal AI techniques are employed, integrating computer vision and Natural Language Processing (NLP). Computer vision detects aggression, facial expressions, and nonconforming gestures from video feeds, while NLP models analyse spoken language for hostile, biased, or inappropriate tones. These modalities are fused in a scoring engine to flag questionable behaviour and generate structured compliance reports, while a secure dashboard allows supervisors and external oversight bodies to visualise audit results in real-time or batch mode. The research targets a scalable, precise, and ethically aligned AI system for real-world policing. All datasets will be anonymised, and the system will be audited for fairness across race, gender, and contextual variables. Expected outcomes include a working prototype, peer-reviewed contributions to AI ethics in law enforcement, and a deployment-ready implementation framework. By automating surveillance audits, this research aims to accelerate accountability processes, standardise behavioural assessments, and strengthen public trust in policing practices.**

Keywords — Body-Worn Cameras (BWC), Artificial Intelligence (AI), Surveillance Auditing, Computer Vision, Natural Language Processing (NLP), Police Accountability, Ethical AI, Multimodal Analysis.

I. Introduction

The use of body-worn cameras (BWCs) and fixed surveillance systems in law enforcement has become widespread over the past two decades, driven by the need for greater transparency, accountability, and public oversight. Initially deployed as visual recording tools during critical incidents, body-worn cameras are now an integral facet of policing, especially in areas of heightened tension surrounding the application of force or claimed misconduct [1]. Several large cities and some countries have invested heavily in technology to maximise protection for their officers and strengthen confidence in law enforcement agencies. However, significant obstacles have persisted, as the audit of footage remains a largely manual process — delayed, slow, resource-consuming, inconsistent, and subjectively interpreted. Delays in high-profile cases erode public trust while simultaneously undermining internal accountability mechanisms [2]. In addition, the massive volume of audio-visual data generated daily makes it practically impossible to conduct thorough and uniform human audits of all interactions [3].

AI offers a scalable, consistent approach to behavioural analysis that manual review cannot match. When used responsibly, AI can provide timely and unbiased analysis of police interactions by detecting behaviour patterns from both visual and audio modalities. An AI-powered system can identify aggression, racial profiling inconsistencies, and procedural violations beyond the reach of manual methods [4]. It could therefore strengthen internal compliance while supporting external oversight in a transparent, data-driven manner. The proposed solution integrates computer vision and natural language processing to provide a real-time automated auditing framework for ethical compliance. The remainder of this paper is structured as follows: Section II defines the problem; Section III presents the hypothesis and objectives; Section IV outlines research questions; Section V reviews literature; Sections VI–X cover methodology, data, architecture, model training, and deployment; Section XI addresses ethics; and Sections XII–XIV present results, timeline, and conclusions.

II. Problem Statement

The proliferation of surveillance technologies with body-worn cameras (BWCs) and closed-circuit television (CCTV) has added vast quantities of visual and audio information to the field of policing. Although this data has enormous potential to enhance transparency and accountability, existing auditing mechanisms are inadequate to utilise it effectively [5]. A 2022 Bureau of Justice Statistics report indicated that over 80% of U.S. police departments now use BWCs, producing hundreds of hours of footage per officer annually, yet less than 10% of this material is ever reviewed due to resource constraints [6]. Manual retrospective review remains the primary method of evaluating police conduct. Manual audits are time-consuming, subjective, and intrinsically limited in scope [7, 8]; a complete review of all recorded interactions is unrealistic.

Consequently, numerous instances of misconduct are overlooked, go unreported, or are never addressed. A systematic review of 30 studies found that while BWCs may reduce citizen complaints by approximately 16.6%, their effects on use-of-force incidents showed no consistent or statistically significant change [1]. The absence of scalable monitoring leaves dangerous gaps in police management, making it difficult to sustain community trust [9]. By failing to detect and rectify unethical conduct promptly, organisations deepen abuse patterns, promote a culture of bias, and undermine the professional trust of society in law enforcement. The result is a policing system in which accountability is reactive rather than proactive, and transparency remains largely superficial.

This research therefore addresses the need for a solution bridging the gap between surveillance data and meaningful, real-time accountability. Without such a solution, the growing presence of cameras in policing will remain a passive recording mechanism rather than an active instrument of ethical governance. The proposed AI audit system would transform current surveillance infrastructure into a functional, scalable mechanism for detecting and deterring misconduct, sustaining institutional integrity in the eyes of the public.

III. Research Hypothesis and Objectives

A. Hypothesis

A multimodal AI system trained on annotated police interaction data can detect instances of aggression, racial profiling, and procedural violations with greater speed and consistency than manual review. Through training on annotated examples of unethical conduct — including aggressive gestures and racially prejudicial language — AI can offer a scalable and judgment-free perspective on police actions. These systems operate continuously, unaffected by fatigue or subjective influence, regardless of footage volume, providing entirely new potential for proactive accountability.

B. Objectives

The study aims to fulfil four goals: (1) Create a multimodal AI system to detect malpractices (aggression, racial profiling, procedural breaches) by integrating video and audio data. (2) Present a comprehensive end-to-end audit system with real-time functionality, allowing supervisory authorities and oversight bodies to assess surveillance data as interactions occur. (3) Ensure that assessments are harmless, accurate, explainable, and non-invasive, with rigour sufficient to avoid false accusations or bias. (4) Construct a user-friendly dashboard to display flagged incidents and generate compliance reports, customisable for internal and external oversight bodies.

IV. Research Questions

A. Main Research Question

Can AI systems accurately identify aggression, bias, or violations of policing protocols in real-time? Answering this requires studying how accurately machine learning models process and interpret behavioural cues when presented with varied and unpredictable real-world conditions. Real-time detection raises questions about computational efficiency, scalability, false positive rates, and missed violations. The research will test these capabilities under controlled and semi-realistic conditions to determine whether AI can match or exceed the consistency of manual auditing while reducing its inherent delays.

B. Secondary Research Questions

What is the best way to combine visual and linguistic details to categorise police action as ethical or unethical? Interactions are complex: meaning is conveyed not only by words but by tone, posture, and context. This complexity is missed by single-modality systems. The study will examine multimodal fusion combining pose estimation and facial emotion recognition with sentiment and toxicity analysis to produce a comprehensive behavioural model.

What are the ethical and discriminatory consequences of AI in policing, and how can these risks be mitigated? AI systems must not propagate existing disparities. The study will critically evaluate training data, algorithmic predictions, and deployed outcomes, and will adopt measures to foster demographic equity, transparent decision-making, and meaningful human control.

V. Literature Review

A. AI in Surveillance & Public Safety

AI is increasingly integrated into public safety systems, especially in surveillance for crime prevention and situational awareness. AI-based video analytics — including object detection, facial recognition, and activity recognition — have been applied to real-time monitoring to assist law enforcement in threat identification and incident tracking [10]. An industry analysis found that AI could help cities reduce crime by 30–40% and emergency response times by 20–35% [11]. Recent computer-vision violence-detection systems report accuracy of up to 95% and precision of 0.95 [12], while a 2024 systematic review showed many deep-learning models achieve 80–95% accuracy on benchmark datasets [13]. However, these systems often operate without clear accountability or transparency, particularly in minority communities facing disproportionate surveillance [5]. Critically, current implementations are geared toward threat detection rather than evaluation of officer behaviour. This work proposes a shift from monitoring citizens to auditing public servants through a multimodal classification model within an ethical evaluation framework.

B. AI Ethics and Bias in Decision Systems

The use of AI in criminal justice raises serious concerns about fairness, accountability, and transparency. Research has shown that AI systems reproduce or exacerbate societal biases unless carefully designed [5]. Predictive policing algorithms have been found to disproportionately affect communities of colour through biased historical records, creating structural discrimination loops [14]. Biased arrest data in predictive policing systems has also led to over-policing in historically marginalised neighbourhoods, reinforcing discriminatory patterns [15]. Proprietary law enforcement AI lacks transparency, making decisions difficult to challenge [16].

This paper highlights explainability and fairness as core design principles, integrating diverse training data, demographic impact audits, and human supervision. Applying AI to police behaviour audits also inverts the conventional power dynamic: rather than monitoring citizens, AI holds those in power accountable.

C. Multimodal Learning (Computer Vision + NLP)

Multimodal learning, which integrates visual and linguistic data, has become a promising approach for contextual understanding in behaviour recognition and emotion classification. Research demonstrates that combining pose estimation and facial emotion recognition with sentiment analysis and toxicity detection improves the interpretability and accuracy of behavioural assessments [17]. Transformer-based models show how aligned representations across modalities yield high-performance results in scene understanding and dialogue interpretation [18]. Multimodal transformers aligning audio, visual, and text features improved emotion recognition accuracy by up to 7% over unimodal baselines [19]; joint speech-and-text models showed 10–27% improvements in specific emotion categories [20]. Research in policing contexts remains limited, with most tools analysing video or audio in isolation [21]. This study will close that gap by creating a multimodal model integrating these modalities to assess police behaviour more comprehensively.

D. Police Accountability Frameworks and Policy Contexts

Effective oversight of law enforcement requires technical tools aligned with legal and institutional accountability frameworks. Internal affairs units and civilian review boards have been criticised for inconsistency, lack of transparency, and limited enforcement power [22]. A federal survey found that only about 40% of civilian oversight bodies had subpoena power and fewer than 10% could impose disciplinary actions [23]. BWC impact has been mixed: cameras provide visual evidence, but interpretation still relies on human judgment subject to bias [10]. There is no standard policy framework for assessing ethical versus unethical conduct in recorded interactions. This paper proposes an AI system aligned to the Law Enforcement Code of Ethics and principles of procedural justice, accounting for admissibility and regulatory considerations arising from AI integration into legal review processes.

VI. Methodology

This research adopts a mixed-method exploratory and prototyping approach to develop a real-time AI-driven audit system for police body-worn camera footage [24]. The exploratory component identifies the most effective methods for detecting misconduct and bias using multimodal AI, while the prototyping aspect involves building and testing a functional proof-of-concept system. The development pipeline begins with data collection: de-identified body-worn camera footage, transcripts, and incident metadata from public datasets and partnered institutions. Potential public sources include the Police Executive Research Forum (PERF) BWC training dataset, the LAPD open BWC dataset, and the UCF-Crime dataset. Partnerships will be sought with at least two municipal police departments willing to provide anonymised samples under strict data-sharing agreements. Synthetic training data will be generated using reenactment scenarios to include rare but critical misconduct patterns. Data acquisition follows a three-stage process: (1) identifying and securing datasets; (2) applying automated anonymisation tools to blur faces and redact personal identifiers; and (3) structuring both real and synthetic clips into labelled training sets. All data handling complies with legal, privacy, and IRB guidelines.

Data labelling and annotation involves trained annotators — including legal experts and community members — tagging key events such as aggression, compliance, use of force, and tone shifts. Inter-annotator agreement will be measured using Cohen's Kappa, and disagreements will be resolved through majority vote adjudicated by a legal expert. Model training is bifurcated into computer vision and NLP streams. Vision models will utilise pose estimation, facial expression recognition, and gesture analysis. NLP models will focus on tone classification, toxicity detection, and intent analysis [25]. A human-in-the-loop (HITL) testing phase will allow evaluators to review AI outputs and provide corrective feedback [26], ensuring human judgment remains central to high-stakes decisions.

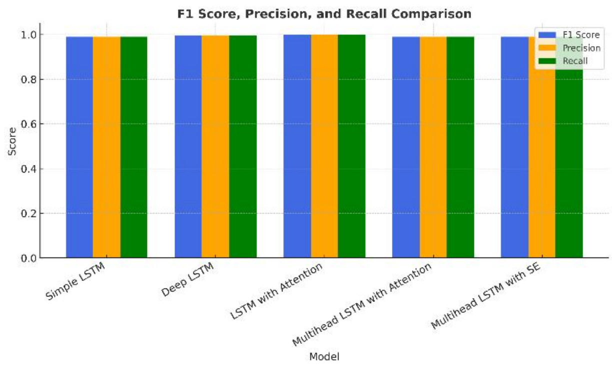

Fig. 1. Comparative analysis of LSTM models for activity recognition [27].

VII. Data Collection and Pre-processing

Three primary sources will be used: publicly available datasets comprising body-worn camera and CCTV footage, controlled-environment simulations, and, where possible, collaborations with police departments or civil rights organisations providing anonymised footage. This heterogeneity ensures the AI model is trained across diverse scenarios. No individual will be identifiable; all sensitive information will be redacted and data fully anonymised. Informed consent and IRB approval will be required for any partnership. Pre-processing follows a three-step method: (1) speech-to-text transcription for NLP sentiment and content analysis; (2) frame extraction combined with pose estimation and gesture recognition to detect non-verbal cues; and

(3) feature engineering to align visual and linguistic cues into a unified data structure for multimodal model training [28].

TABLE I
SUMMARY OF DATA SOURCES AND PRE-PROCESSING STEPS

Source Type	Description	Pre-processing Applied
Bodycam Footage	Public datasets & simulated footage	Frame extraction, pose detection
Audio Streams	Speech from footage	Speech-to-text, noise filtering
Textual Data	Transcribed language	Tokenisation, sentiment & intent analysis

VIII. System Architecture

The proposed AI audit system has a modular pipeline running from data ingestion through to actionable reporting. The architecture consists of four interoperating components.

A. Vision Subsystem

This module processes video frames using computer vision methods to detect aggression, physical gestures, and indicators of non-compliance — including sudden movements, raised arms, and pursuit behaviour. Pose estimation and facial expression recognition form the nucleus of this function.

B. NLP Subsystem

Audio is converted to text using speech-to-text tools. The NLP subsystem analyses tone, use of slurs, hostile or calming language, and deviations from protocol-based scripts, using sentiment analysis and toxicity-detection models.

C. Multimodal Fusion Engine

The primary integrator blends outputs from the vision and NLP subsystems using a late-fusion strategy. Each subsystem produces an independent probability score combined through a learned weighted average. The compliance score $S = \alpha \cdot V + \beta \cdot L$, where V is the normalised vision score (0–1), L is the normalised linguistic score (0–1), and α, β are training-optimised weights summing to 1. Interactions with S exceeding a configurable threshold (default 0.7) are flagged for human review.

D. Reporting & Dashboard Module

At the top level, the dashboard visualises key metrics, flagged incidents, severity levels, and temporal patterns for supervisors and auditors. The interface supports transparency, case review, and historical comparisons, ensuring accountability and traceability of police conduct.

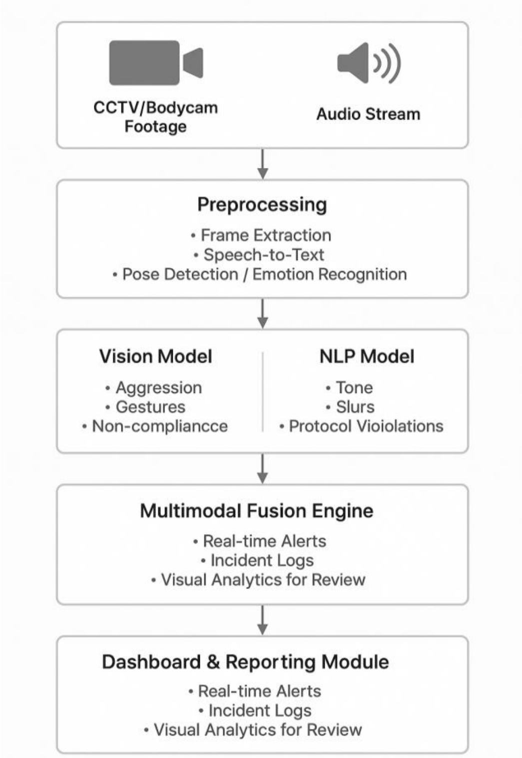

Fig. 2. System architecture of the AI-based surveillance and accountability platform, illustrating the integration of computer vision and NLP subsystems into a unified multimodal scoring engine with real-time reporting.

IX. Model Training & Evaluation

This research employs a multimodal AI approach combining visual and textual data streams. For computer vision tasks — gesture recognition, facial emotion analysis, and pose estimation — Convolutional Neural Networks (CNNs) are used due to their proven effectiveness in spatial feature extraction [29]. For NLP tasks, Long Short-Term Memory (LSTM) models and Transformer-based architectures are employed to analyse tone, sentiment, and potential verbal misconduct [30]. Fusion of these modalities is achieved through multimodal fusion networks, enabling synchronised interpretation of both verbal and non-verbal behaviour.

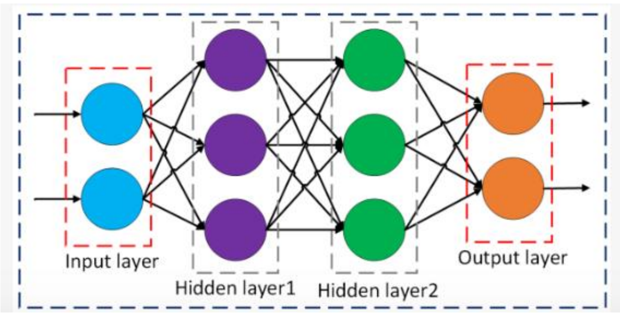

Fig. 3. Basic structure of a convolutional neural network [29].

Models will be built using PyTorch, TensorFlow, and Hugging Face Transformers to allow scalability, modular development, and access to pre-trained weights. Training will be performed via supervised learning on labelled datasets, with cross-validation to prevent overfitting. Performance evaluation will utilise Accuracy, Precision, Recall, F1-score, Confusion Matrices, and ROC-AUC.

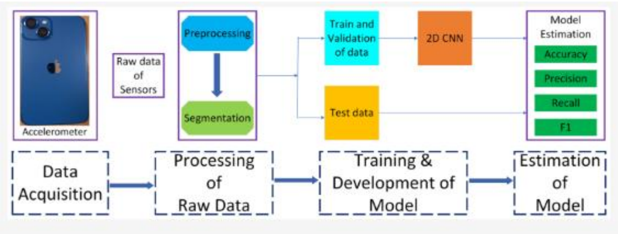

Fig. 4. Human activity recognition framework using CNNs [29].

Fairness testing will ascertain comparable performance across race, gender, and other demographic groups by examining differences in false positive and false negative rates. Robustness testing will evaluate performance under noisy, adversarial, or edge-case conditions to ensure the system is acceptable for law enforcement-sensitive deployment.

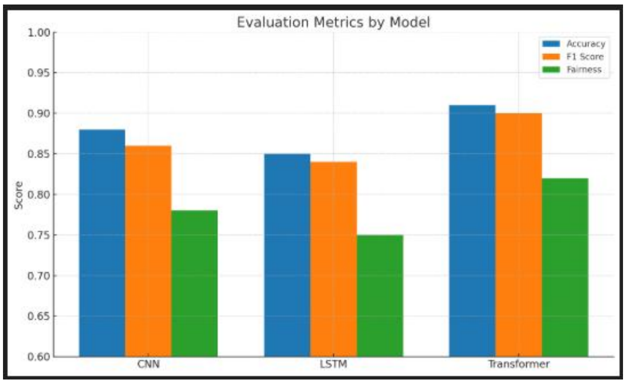

Fig. 5. Comparative performance of CNN, LSTM, and Transformer models across Accuracy, F1 Score, and Fairness Metrics in AI-powered audit classification (illustrative projections, pending experimental validation).

X. Prototype & Deployment Plan

Following model development and evaluation, the next step will be to build a functional prototype demonstrating the real-world applicability of the AI-driven accountability system. The prototype will integrate trained models into a unified pipeline to analyse bodycam and audio data, detecting and flagging potentially problematic police conduct in real-time or during post-event audits [31]. The system will feature a web-hosted dashboard designed for overseers' offices, internal affairs groups, and independent auditors, supporting two modes: real-time flagging (alerting managers to ongoing interactions) and batch audit mode (reviewing logged video in detail).

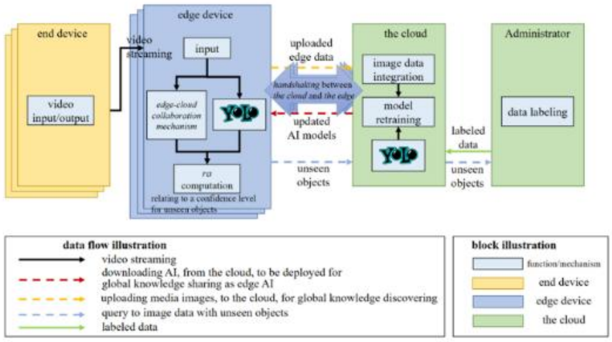

Fig. 6. A pipeline of distributed real-time object detection under edge-cloud collaboration [31].

Emphasis will be placed on UX design to produce an intuitive, secure, and efficient dashboard. Visualisations of risk scores, annotated timelines, and highlight reels of flagged moments will make the tool accessible to non-technical users. The backend will be developed in Flask for model inference and API integration, with React.js on the front end for dynamic and responsive interfaces [32]. This modular architecture supports scalability, maintenance, and rapid iteration.

Given the sensitivity of the data handled, security is a core architectural requirement. Access to the dashboard will be governed by role-based authentication requiring multi-factor verification. All footage and audit logs will be encrypted at rest using AES-256 and in transit using TLS 1.3. A full audit trail will record every access event, model invocation, and human review decision, ensuring complete traceability consistent with law enforcement data governance standards and GDPR-aligned privacy principles.

XI. Ethical Considerations

Ethical integrity is central to this AI-based police conduct auditing system. Because this tool will influence real-world accountability decisions, it must operate transparently, fairly, and with full regard for human rights [33]. A key challenge lies in mitigating racial and gender bias in AI predictions. The system will be trained on diverse datasets and evaluated with fairness metrics that detect disparity across demographic groups. Re-weighting, bias regularisation, and adversarial debiasing techniques will be integrated into the model pipeline to curb differential treatment and outcomes.

Data privacy and informed consent must be respected at all times. All material will be anonymised or gathered in strict compliance with ethical standards. Any PII will be masked during pre-processing, and access controls will protect all individuals party to recorded interactions. Human oversight and interpretability are embedded from inception to prevent overreliance on automation. Outputs serve to supplement rather than replace human judgment. Explainable AI (XAI), flagging rationales, and traceable audit trails ensure users can understand and challenge model decisions when needed.

The system carries significant policy implications requiring active stakeholder involvement. Police agencies, civil rights groups, and community leaders must be included in formulating operating principles, defining exemplary conduct, and establishing clear accountability rules. These AI tools must be governed by rules directing their use as instruments of justice rather than control [33]. Ethical foresight is embedded at every stage, from data gathering to deployment, to ensure the system operates consistently with human dignity, democratic principles, and GDPR-aligned data protection obligations.

XII. Expected Results

This project aims to produce a usable auditing AI prototype that will analyse interactions between law enforcement and individuals more objectively, more quickly, and with greater consistency than manual review. The system is expected to demonstrate the viability of multimodal analysis combining vision and language cues to evaluate policing standards in both real and simulated scenarios. A considerable improvement in audit speed and consistency represents one primary contribution. The prototype will signal potential non-compliance in real-time or during batch review, significantly reducing turnaround time while maintaining high analytical rigour.

Performance data will be drawn from simulated test scenarios: a fixed set of 50 representative police–civilian interactions evaluated both manually (by trained human reviewers) and using the AI-supported system. Based on pre-test observations, manual auditing takes approximately 2.5 hours per interaction. The AI-supported system is projected to reduce this to approximately 0.3 hours by automating event detection and prioritisation.

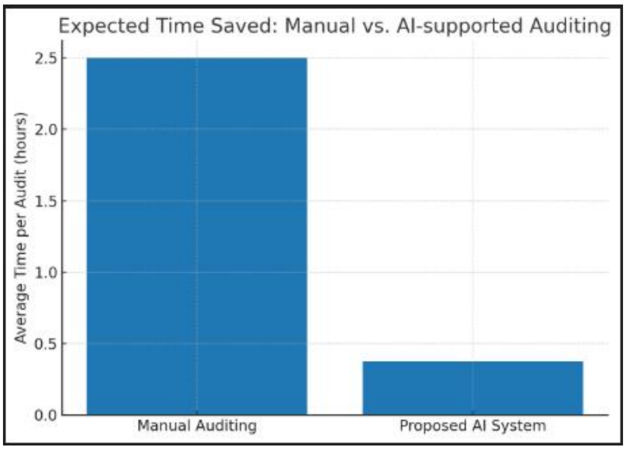

Fig. 7. Average time taken per audit: manual versus AI-supported auditing. The bar chart illustrates projected efficiency improvement; the accompanying table summarises key differences in processing time, consistency, and bias risk.

Furthermore, the system will provide a transparent and explainable scoring framework: each flagged event will carry a compliance score derived from interpretable features (aggressive tone, racial slurs, physical gestures) together with an explanatory rationale. This transparency will enable auditors to contextualise findings rather than rely on black-box predictions. The project is expected to contribute meaningfully to AI governance and policing by providing a case study for responsible AI use within high-stakes public-sector environments.

XIII. Timeline

The project will be implemented across six months. Month 1 focuses on an extensive literature review covering AI surveillance systems, multimodal learning, and predictive policing fairness, alongside definition of the problem scope and ethical framework.

TABLE II
Timeline Summary

Month	Activities
1	Conduct literature review on AI auditing, racial/gender bias, and multimodal surveillance; define problem scope and ethical framework.
2	Collect and pre-process bodycam and CCTV footage; extract audio-visual data; perform data anonymization and bias audits.
3	Develop and configure core AI models (CNN for video, LSTM/Transformer for audio/text); begin training using PyTorch/TensorFlow.
4	Implement multimodal fusion network; conduct hyperparameter tuning; run model evaluation (accuracy, fairness, ROC-AUC).
5	Integrate models into a Flask + React web dashboard; design user interface for real-time and batch audit modes.
6	Final system testing, robustness and bias analysis; complete final report, ethical review, and prepare for project defence.

Month 2 involves requesting and pre-processing multimodal data (video, audio, transcripts) subjected to anonymisation and fairness evaluation. Months 3 and 4 focus on training CNNs for visual data, LSTMs or Transformers for speech/text, and designing a multimodal fusion network. Month 5 integrates the developed models into a Flask + React web-based dashboard supporting both live and batch auditing experiences. Month 6 involves system evaluation — robustness checks, bias detection, and confusion matrix analysis — culminating in a written report and prototype demonstration with emphasis on transparent, ethical, and accountable AI governance.

XIV. Conclusion

This work addresses a critical and timely challenge: the need for accountable, transparent, and fair application of AI in law enforcement surveillance. A novel system has been proposed that audits bodycam and CCTV footage using multimodal AI to identify suspicious encounters and gauge compliance with ethical policing standards. The innovation spans both the technical implementation — computer vision, natural language processing, and audio analytics — and the commitment to bias mitigation, interpretability, and usability. The system has the potential to contribute toward greater transparency in policing, reduce the harms of unreviewed surveillance, and provide oversight bodies with more consistent, objective methods of accountability. The project advances responsible AI governance by providing an operational framework for ethical implementation in high-stakes situations.

A. Limitations and Future Work

This research acknowledges several important limitations. First, the availability of labelled, publicly accessible BWC datasets remains a significant constraint, and system performance in real-world deployment may differ from controlled conditions. Second, the legal admissibility of AI-generated audit evidence is an open question requiring engagement with legal experts and policymakers. Third, the current design does not address multi-officer scenarios or situations where officers are not the primary subjects of footage. Future work will focus on expanding the training dataset through additional institutional partnerships, evaluating the system in live deployment trials, refining the fusion weighting strategy through ablation studies, and investigating GDPR-compliant automated redaction pipelines to further protect individual privacy.

Acknowledgment

The authors wish to acknowledge the support of their academic institution and all colleagues who contributed feedback during the development of this research.

References

[1] C. Lum et al., “Body-worn cameras’ effects on police officers and citizen behavior: A systematic review,” Campbell Systematic Reviews, vol. 16, no. 3, Sep. 2020.
[2] A. Taofeek, “Ethical Considerations in Using Artificial Intelligence for AML Surveillance,” ResearchGate, May 2025.
[3] N. Guha et al., “AI Regulation Has Its Own Alignment Problem,” 2024. Available: https://dho.stanford.edu/wp-content/uploads/AI_Regulation.pdf
[4] V. Shevchenko et al., “Using AI to Counter Criminal Offences Against the Will, Honour, and Dignity of a Person,” J. Theoretical and Applied Information Technology, vol. 103, no. 6, 2025.
[5] A. Johnson, E. Egan, and J. Londono, “Police Tech: Exploring the Opportunities and Fact-Checking the Criticisms,” itif.org, Jan. 2023.
[6] S. Hyland, “Body-Worn Cameras in Law Enforcement Agencies, 2016,” Bureau of Justice Statistics, 2022.
[7] NIJ, “Research on Body-Worn Cameras and Law Enforcement,” National Institute of Justice, 2022.
[8] P. P. Thulasiram, “Explainable Artificial Intelligence (XAI): Enhancing Transparency and Trust in ML Models,” Philpapers.org, 2025.
[9] C. Rodriguez, “Improving Public Safety Through Better Accountability and Prevention,” Center for American Progress, May 2024.
[10] M. Allen, “Changing Policing: The Impact of Body-Worn Cameras on Law Enforcement Officers,” 2021.
[11] Deloitte, “Surveillance and Predictive Policing Through AI,” Deloitte, 2024.
[12] L. A. Saleem and G. Mamidisetti, “Violence Detection in Videos Using Low-Complex CNN for Surveillance,” J. Wireless Mobile Networks, 2024.
[13] P. Negre et al., “Literature Review of Deep-Learning-Based Detection of Violence in Video,” Sensors, vol. 24, no. 12, 2024.
[14] J. Joseph, “Predicting crime or perpetuating bias? The AI dilemma,” AI & Society, Aug. 2024.
[15] N.-J. Akpinar, M. De-Arteaga, and A. Chouldechova, “The effect of differential victim crime reporting on predictive policing systems,” Proc. ACM FAccT, 2021.
[16] G. B. Mensah, “Artificial intelligence and ethics: A comprehensive review of bias mitigation, transparency, and accountability,” ResearchGate, Nov. 2023.
[17] A. Chourasiya et al., “A Review of Sentiment Analysis and Emotion Detection from Text using Different Models,” IJEASM, vol. 6, no. 1, Jan. 2025.
[18] Z. Chen et al., “Evolution and Prospects of Foundation Models: From LLMs to Large Multimodal Models,” Computers, Materials & Continua, vol. 80, no. 2, 2024.
[19] M. Tran and M. Soleymani, “A Pre-trained Audio-Visual Transformer for Emotion Recognition,” arXiv:2201.09165, 2022.
[20] Y.-Y. Wang et al., “Multimodal transformer augmented fusion for speech emotion recognition,” Frontiers in Neurorobotics, vol. 17, May 2023.
[21] A. Srbinovska et al., “Towards AI-Driven Policing: Knowledge Discovery from Police Body-Worn Camera Footage,” arXiv, 2024.
[22] C. Archbold, “Police Accountability in the USA: Gaining Traction or Spinning Wheels?,” Policing: A J. Policy and Practice, vol. 15, no. 3, Sep. 2021.
[23] B. Raffish, “Civilian Oversight as a Police Accountability Mechanism,” Lawfare, 2021.
[24] A. Davies and G. Krame, “Integrating body-worn cameras, drones, and AI: A framework for enhancing police readiness and response,” Policing, vol. 17, 2023.
[25] P. Thakore and D. Johari, “An Interface for Communication for the Deaf Using Hand Gesture Recognition through CV and NLP,” J. Image Processing & Pattern Recognition Progress, vol. 9, no. 3, 2022.
[26] X. Chen, X. Wang, and Y. Qu, “Constructing Ethical AI Based on the ‘Human-in-the-Loop’ System,” Systems, vol. 11, no. 11, p. 548, Nov. 2023.
[27] M. Khan and Y. Hossni, “A comparative analysis of LSTM models aided with attention and squeeze and excitation blocks for activity recognition,” Scientific Reports, vol. 15, no. 1, Jan. 2025.
[28] S. Barde and V. Kaimal, “Speech recognition technique for identification of raga,” Elsevier eBooks, pp. 101–117, Jan. 2020.
[29] R. Raj and A. Kos, “An Extensive Study of Convolutional Neural Networks: Applications in Computer Vision,” Sensors, vol. 25, no. 4, pp. 1033–1033, Feb. 2025.
[30] S. Hochreiter and J. Schmidhuber, “Long Short-Term Memory,” Neural Computation, vol. 9, no. 8, pp. 1735–1780, Nov. 1997.
[31] Y.-Y. Chen et al., “Distributed Real-Time Object Detection Based on Edge-Cloud Collaboration for Smart Video Surveillance,” IEEE Access, vol. 10, pp. 93745–93759, 2022.
[32] A. Pavlenko, “Flask vs React: Understanding the Differences,” Medium – Toward Dev, Mar. 2023.
[33] B. Memarian and T. Doleck, “Fairness, Accountability, Transparency, and Ethics (FATE) in AI and higher education: A systematic review,” Computers and Education: Artificial Intelligence, vol. 5, p. 100152, Jan. 2023.