Accident Detection System

Suhas H [1], Assit.Prof. Vidya S [2]

[1] Student, Department of MCA, Bangalore Institute of Technology, Karnataka, India

[2] Assistant Professor, Department of MCA, Bangalore Institute of Technology, Karnataka, India

Abstract

Road traffic accidents remain a major global public-health challenge, where delays in detection and poor contextual information often increase the severity of outcomes. This paper presents ADS AI, a practical, end-to-end framework that combines a lightweight, low-latency detection stage with a deeper, multi-modal reasoning stage to deliver fast, actionable intelligence for emergency response. The first tier employs a YOLOv8 object-detection model for near-real-time identification of anomalous events from live video streams; the second tier uses a multi-modal generative model to create structured, human-readable incident summaries and recommendations. The platform integrates persistent storage, geospatial analytics, and automated alerting via messaging APIs to form a deployable prototype. We describe the architecture, implementation choices, evaluation approach, and future directions, highlighting how this two-tiered design enables both speed and situational depth—shifting traffic safety systems from reactive monitoring toward data-driven, proactive operations.

Keywords: accident detection, YOLOv8, generative AI, real-time surveillance, emergency alerting, geospatial analytics, Streamlit

1. Introduction

Urbanization, rising vehicle numbers, and increasingly complex traffic patterns have made rapid detection and context-aware response to road accidents more important than ever. Conventional CCTV monitoring and manual reporting pipelines introduce substantial delays and often lack the contextual detail needed by first responders. Modern computer vision offers high-speed perception while recent advances in multi-modal generative AI provide the ability to synthesize concise scene descriptions and suggested actions. ADS AI combines these strengths in a modular, practical system designed to reduce time-to-action and deliver richer incident information to emergency services and planners.

This paper documents the design and prototype implementation of ADS AI, explains the engineering trade-offs that guided the architecture, and reports on the system's core capabilities. The contribution is a reproducible blueprint for integrating fast, single-stage detectors with slower, reasoning-oriented models to produce an operationally useful accident-management pipeline.

2. Related Work

Research in automated traffic incident management spans several strands:

- Classical computer vision and handcrafted features: Early work used background subtraction, optical flow and descriptors such as HOG/LBP with traditional classifiers to identify anomalies. These methods are interpretable but fragile in unconstrained environments.
- Deep learning for detection and temporal analysis: CNNs and hybrid CNN-RNN models improved robustness by learning hierarchical representations. Single-stage detectors in the YOLO family balance accuracy and latency, making them attractive for live surveillance.
- Generative and large multi-modal models: Emerging research explores using LLMs and multi-modal models to generate narratives, label complex scenes, and recommend response actions—extending detection into comprehension.
- End-to-end systems and alerting: Several efforts combine sensing, storage, and messaging systems (including Telegram bots and GIS tools) to close the loop between detection and action. ADS AI brings these elements together into a two-tier AI pipeline designed for low latency and rich context.

3. Problem Statement

Existing automated solutions frequently stop at detection: an alarm is raised but decision-relevant context (vehicle types, likely injuries, environmental factors, nearby medical resources) is missing. Emergency services therefore operate with incomplete situational awareness. The challenge is to create a system that (1) detects incidents with near real-time responsiveness, (2) produces structured, actionable contextual summaries, and (3) persists data for analytics and planning—while remaining feasible to deploy in constrained environments.

4. Proposed System

ADS AI adopts a two-tier AI architecture:

Tier 1 — Fast Perception (YOLOv8): A YOLOv8 model processes video frames to detect vehicles, pedestrians, and anomalous motion patterns. Detection results are subject to temporal confirmation logic to reduce spurious alarms (e.g., require detection over N consecutive frames). The objective is sub-second detection latency suitable for live feeds.

Tier 2 — Deep Scene Comprehension (Generative AI): Once an incident is confirmed, a representative keyframe is forwarded to a multi-modal generative model (e.g., Google Gemini or equivalent). A carefully engineered prompt requests a structured output: a concise scenario description, assessed severity, likely vehicle types involved, environmental observations, potential causes, recommended first-response actions, and suggested resources (nearest hospitals, traffic diversions).

Supporting subsystems: A NoSQL database stores incident records and AI reports. Geospatial services resolve coordinates to addresses and enable hotspot mapping. A messaging client (Telegram Bot API) dispatches media-rich alerts to configured responders. A lightweight analytics dashboard (Streamlit) exposes live monitoring and historical insights.

5. Methodology

The system accepts live video streams or uploaded recordings. Training and evaluation use annotated datasets containing accident and non-accident events, enriched with metadata where available (time-of-day, weather, road type). Preprocessing includes frame extraction, normalization, and augmentation to improve generalization across lighting and weather conditions.

5.2 Model design and training

- Detection: We adopt a YOLOv8 backbone pre-trained on large detection datasets and fine-tune it on domain-specific accident imagery. Hyperparameter search focuses on detection confidence thresholds and IoU settings appropriate for small, occluded objects.

- Severity heuristics: Lightweight rule-based heuristics operating on detection dynamics (e.g., intersecting bounding boxes, sudden changes in velocity, deformation proxy via overlapping areas) produce an initial severity label (Low / Moderate / High) to guide responder prioritization.

- Generative prompts: Prompt templates are engineered to solicit structured JSON-like outputs from the generative model, ensuring predictable fields and terse recommendations suitable for rapid consumption by humans and machines.

5.3 Evaluation metrics

Detection performance is measured with precision/recall and mAP at multiple IoU thresholds. System-level evaluation includes detection latency, false alarm rate (after confirmation), and qualitative assessment of generative summaries by domain experts (usefulness, accuracy, and actionability).

6. System Architecture and Implementation

ADS AI is implemented as a modular web application with three conceptual layers:

- Presentation Layer (Streamlit): Live video display, parameter controls, and dashboards.

- Application Layer: Video ingestion, YOLO inference, confirmation logic, AI orchestration (callouts to the generative model), severity heuristics, and alerting logic.

- Data & Integration Layer: MongoDB for flexible incident storage, filesystem or object store for media, geocoding services for reverse lookup, and the Telegram API for alerts.

Technology highlights: Python 3.x, Ultralytics YOLOv8, PyTorch, OpenCV, Streamlit for lightweight dashboards, Plotly/Folium for visualizations, MongoDB for persistence, and HTTP clients for external APIs.

Reliability measures: Model loading is cached to avoid repeated I/O; frame confirmation reduces transient false positives; alert throttling and rate limits prevent spamming responders.

7. Prototype Functionality and Results

A working prototype demonstrates the following capabilities:

- Near-real-time detection: YOLOv8 identifies vehicles and anomalous events with low latency. A configurable frame-confirmation window reduces false positives in noisy scenes.

- Automated severity tagging: Heuristics give an immediate severity estimate usable for triage.

- Generative incident summaries: The multi-modal model returns structured scene descriptions and recommended actions that complement sensor data for responder decision-making.

- Media-rich alerting: The system can dispatch images, coordinates, and the AI-generated summary to a preconfigured Telegram channel, including reverse-geocoded addresses.
- Analytics: Stored incident records enable heatmap generation and hotspot identification to support proactive infrastructure interventions.

Quantitative results will vary with the dataset and deployment environment; the prototype shows promising detection precision with practical end-to-end latencies for a single camera instance.

8. Discussion, Limitations and Future Work

Limitations: The model's robustness depends on training data diversity—region- and camera-specific fine-tuning are recommended. Dependence on third-party cloud services introduces availability and privacy considerations. Scaling to many parallel streams requires containerized architectures and orchestration.

Future directions: - Cloud & edge hybridization: Deploy lightweight detectors at the edge for low latency, forwarding only confirmed events to central servers for deeper analysis. - MLOps for continuous learning: Implement pipelines to collect labeled false positives/negatives and retrain models safely in production. - Privacy-preserving modes: Add blurring and redaction steps to comply with local regulations and protect personal data. - Expanded multi-modal fusion: Integrate non-visual telemetry (e.g., vehicle telemetry, roadway sensors) to improve inference and early warning.

9. Conclusion

ADS AI illustrates how pairing a fast perceptual model with a slower, reasoning-oriented generative model can produce a practical system that both detects incidents rapidly and provides rich contextual intelligence to support emergency response and planning. The resulting prototype is a reproducible template for municipalities and researchers seeking to move from passive surveillance to proactive safety management.

References

[1]. K. Wu, W. Li, and X. Xiao, "AccidentGPT: Large Multi-Modal Foundation Model for Traffic Accident Analysis," arXiv preprint arXiv:2401.03040, 2024.
[2]. H. Ghahremannezhad, H. Shi, and C. Liu, "Real-Time Accident Detection in Traffic Surveillance Using Deep Learning," arXiv preprint arXiv:2208.06461, 2022.
[3]. "Real-Time Vehicle Accident Recognition from Traffic Video Surveillance using YOLOv8," International Journal on Recent and Innovation Trends in Computing and Communication, vol. 11, no. 5, 2023.
[4]. "Accident Detection System Using YOLOv8 and CNN," International Research Journal of Modernization in Engineering, Technology and Science, vol. 5, no. 12, 2023.
[5]. "SeeUnsafe: A New Framework for Video-Based Traffic Accident Analysis," arXiv preprint arXiv:2501.10604, 2025.
[6]. "LLM Multimodal Traffic Accident Forecasting," PubMed Abstract 38005612, Nov. 2023.
[7]. M. S. R. Rajesh and M. Das, "Detection of Road Accidents using Computer Vision with Generative AI," International Journal of Research Publication and Reviews, vol. 5, no. 11, 2024.
[8]. "IoT Cloud System for Traffic Monitoring and Vehicular Accidents," IFTM University Research Paper, 2023.
[9]. "IoT Based Accident and Analysis System Using Edge Computing," International Journal of Information and Research in Engineering, 2024.
[10]. "IoT Enabled Accident Detection and Rescue System," International Research Journal of Modernization in Engineering, vol. 7, no. 3, 2025.
[11]. "Smart Accidental Alert and Location Tracking System Using Telegram," International Research Journal of Engineering and Technology, vol. 9, no. 12, 2022.
[12]. "Development of Accident Detection System Using YOLOv8 with Telegram Notification," International Journal of Computer Science, vol. 13, no. 4, 2024.
[13]. "Real Time Accident Detection and Ambulance Rescue using Deep Learning," International Journal of Advanced Research in Science, Communication and Technology, vol. 4, no. 2, 2024.
[14]. "Accident Detection Model using YOLOv8 and Roboflow," Roboflow Universe Project, 2024.
[15]. "Road Accident Detection-YOLOv8 Dataset," Roboflow Universe Dataset, 2024.
[16]. "Vehicle Accident Detection Computer Vision Project," Roboflow Universe Project, 2025.
[17]. "Real-Time Traffic Monitoring using Computer Vision and Deep Learning," International Research

Journal of Education and Technology, vol. 6, no. 9, 2024.

[18]. K. K. Jyothi, G. Kalyani, and G. J. Lakshmi, "Computer Vision Based Accident Detection and Alert System," International Journal for Research Trends and Innovation, 2023.

[19]. "Safe Road AI: Real-Time Accident Detection from Multi-Angle Videos," International Research Journal of Engineering and Technology, vol. 11, no. 9, 2024.

[20]. "Deep Learning Ensemble Approach for Traffic Accident Detection," Journal of Engineering Sciences, vol. 16, no. 4, 2025.

[21]. "Road Accident Detection using SVM and Ensemble Learning," International Journal of Advanced Computer Science and Applications, vol. 15, no. 5, 2024.

[22]. "Car Accident Detection Using YOLOv8 Algorithms," International Journal of Innovative Research in Technology, 2024.

[23]. "Enhancing Vehicle Safety: Comprehensive Accident Detection and Alert System," International Journal of Advanced Computer Science and Applications, vol. 14, no. 11, 2023.

[24]. "Traffic Pattern Recognition Using IoT Sensors and Machine Learning," International Journal of Management & Information Systems, vol. 9, no. 1, 2024.

[25]. "Accident Detection and Alert System with Real-Time Capabilities," Science Excel Publication, 2024.