AI-BASED REAL-TIME ROAD ACCIDENT DETECTION AND EMERGENCY RESPONSE SYSTEM USING COMPUTER VISION AND IoT: A LITERATURE REVIEW

Prasad Toshniwal | Pankaj Jadhav | Namrata Rajbhar | Seema Rode

Master of Computer Applications (MCA)
D. Y. Patil Institute of Master of Computer Application and Management, Akurdi, Pune, Maharashtra, India

Abstract

Road traffic accidents require rapid detection and reliable communication of incident information to emergency responders. Conventional reporting often depends on witnesses or manual observation, while automated solutions based only on video or vehicle sensors can fail under difficult conditions or generate false alerts. This literature review examines research on computer-vision-based accident detection, deep-learning object detection, vehicle tracking, Internet of Things (IoT) sensing, automatic license plate recognition (ALPR), and lightweight communication. Studies on traffic-video anomaly detection show that spatiotemporal models can identify unusual events, while trajectory-based methods can detect conflicts among road users. ALPR surveys identify illumination, motion blur, occlusion, camera position, and regional plate variation as persistent challenges. These findings support a multimodal approach in which video analytics are combined with impact and motion sensors, GPS coordinates, vehicle identification, and a communication layer such as MQTT. Edge processing may reduce dependence on continuous cloud round trips, while cloud services can support centralized records and monitoring. However, the reviewed literature does not establish that one integrated pipeline will automatically improve every metric. Further evaluation is required for false-positive rate, detection latency, severity classification, plate-recognition performance, communication reliability, privacy, and emergency-service integration.

Keywords

Road accident detection; computer vision; Internet of Things; YOLO; sensor fusion; emergency response; vehicle tracking; MQTT.

1. Introduction

Road traffic incidents can develop faster than a human observer can report them, particularly on highways, poorly monitored roads, and during night-time conditions. The time between an accident and the delivery of useful information to emergency responders is therefore an important system-design concern. Manual reporting may be delayed when drivers or passengers are injured, unconscious, or unable to communicate. Automated monitoring offers a way to identify potential incidents from camera streams and sensor measurements, but a detection output is useful only when it is sufficiently reliable and accompanied by actionable information.

Computer vision can analyze traffic video to detect vehicles, track their movement, and identify unusual trajectories or collision-like events. Deep-learning detectors such as the YOLO family are designed for object detection, while tracking algorithms associate detections across frames to estimate movement over time [3], [4], [6]–[8]. Accident detection is more complex than vehicle detection alone: a system must infer whether observed motion represents a collision or another event such as sudden braking, congestion, occlusion, or a near miss.

IoT devices can provide complementary evidence. Accelerometers and gyroscopes can measure impact-like acceleration and changes in orientation, while GPS can supply approximate incident coordinates. Automatic license plate recognition (ALPR) and optical character recognition (OCR) may help identify a vehicle, subject to image quality and privacy constraints. MQTT is a lightweight publish/subscribe messaging protocol standardized for machine-to-machine and IoT communications [5]. A practical architecture can use edge devices for low-latency initial analysis and cloud infrastructure for persistent records, dashboards, and further analytics.

The purpose of this literature review is to compare research and technical foundations relevant to an integrated accident-to-response pipeline. The review focuses on what existing work supports, where limitations remain, and how a combined design could be evaluated. It does not claim that the proposed project has already been implemented or experimentally validated. Recent surveys of vision-based traffic accident detection and anticipation emphasize the difficulty of rare, unpredictable events and the need for robust evaluation across datasets [11]. Reviews of IoT-based notification systems similarly identify sensor reliability, communication, and interoperability as continuing challenges [12].

2. Literature Review

2.1 Review Approach and Scope

This paper uses a focused narrative literature review rather than a systematic review or a new experimental study. The selected sources cover video-based accident detection, object detection, multi-object tracking, license-plate recognition, IoT messaging, and recent surveys of accident detection and prediction. Sources were selected for direct relevance to one or more stages of an accident-to-alert workflow. The comparison summarizes the reported purpose and known limitations of each source; performance values from different

papers are not compared as if they came from a common dataset. The proposed architecture is a synthesis for future implementation and has not been experimentally validated in this paper.

2.2 Summary of Existing Research

Pawar and Attar formulate traffic accident detection as a video anomaly-detection problem and investigate spatiotemporal autoencoders and sequence-to-sequence long short-term memory autoencoders [1]. Their work is relevant because accident-related behavior is temporal: a single frame may show two vehicles close together without revealing whether a collision occurred. Anomaly-based systems can reduce the need to enumerate every accident type, but their performance depends on training data, scene characteristics, and the definition of normal behavior.

Ghahremannezhad, Shi, and Liu propose a real-time traffic-surveillance framework that combines YOLOv4 object detection, object tracking, and trajectory-conflict analysis [2]. The method evaluates movement patterns to identify conflicts involving vehicles, pedestrians, and bicycles. This illustrates why object detection should be followed by temporal reasoning rather than treated as a complete accident classifier. Trajectory-based approaches may nevertheless struggle when objects are heavily occluded, tracking identities are lost, or camera geometry is unsuitable.

YOLO methods are widely used for efficient object detection. Redmon et al. introduced the original You Only Look Once approach as a unified, single-pass detection framework [3]. Bochkovskiy et al. later introduced YOLOv4 with training and architectural improvements [4]. Ultralytics provides YOLOv8 as software and documentation rather than a formal peer-reviewed paper, so it should be cited as software when used in implementation [6]. For this project, YOLOv8 can be considered a candidate detector, but its suitability must be measured on the selected road-scene dataset and target hardware.

Vehicle tracking is another necessary component. Bewley et al. introduced SORT, a simple online tracking method that combines Kalman filtering with data association [7]. Wojke, Bewley, and Paulus introduced Deep SORT, adding appearance information to improve identity association through occlusion [8]. These methods are general multi-object trackers rather than accident detectors. Their relevance is that tracking can provide trajectories and temporal context for downstream collision analysis.

Vehicle identification presents a separate challenge. Shashirangana et al. review ALPR methods and identify issues such as changing illumination, motion blur, occlusion, viewing angle, plate-format differences, and computational cost [9]. Khan et al. also survey automatic number-plate recognition approaches for intelligent transportation applications [10]. These reviews suggest that plate recognition should be treated as a confidence-bearing module rather than a guaranteed output. If plate confidence is low, the system should send the incident location and available scene evidence without asserting an uncertain plate number.

The OASIS MQTT Version 5.0 specification defines a lightweight publish/subscribe messaging protocol suitable for constrained devices and IoT communication [5]. MQTT can transmit structured incident messages between edge devices, brokers, and monitoring services, but it does not itself verify that an accident occurred or guarantee that an emergency authority will act on a notification. Application-level authentication, delivery handling, retries, timestamps, and acknowledgement monitoring are still required.

Across these studies, an important distinction emerges between detecting an abnormal event and completing an emergency response workflow. Video analytics provide contextual evidence, sensor systems can provide physical measurements, GPS provides location, ALPR can contribute vehicle identity, and messaging protocols support information transfer. These capabilities are complementary, but each has its own uncertainty and failure modes. A robust system must fuse evidence, preserve confidence information, and provide fallback behavior when one component is unavailable. Broader reviews of traffic-safety deep learning and non-visual accident prediction also show that visual and sensor-based approaches address related but different evidence sources, so their evaluation should not be treated as interchangeable [13], [14].

2.2 Literature Review Comparison Table

No.	Author / Year	Research topic	Method / technology	Main finding and limitation
1	Pawar & Attar, 2022 [1]	Accident detection and localization	Spatiotemporal autoencoder; sequence-to-sequence LSTM autoencoder	Models unusual temporal patterns in traffic video. Scene and dataset variation can affect generalization; not a complete emergency pipeline.
2	Ghahremannezhad et al., 2022 [2]	Real-time accident detection at intersections	YOLOv4, Kalman/Hungarian tracking, trajectory-conflict analysis	Combines object detection and motion analysis. Occlusion, camera geometry, tracking errors, and event ambiguity remain concerns.

3	Redmon et al., 2016 [3]	General real-time object detection	YOLO single-stage detector	Introduces unified detection with efficient inference; not designed specifically to classify road accidents.
4	Bochkovskiy et al., 2020 [4]	Efficient object detection	YOLOv4 architecture and training strategies	Improves the detection framework. Performance depends on data, model, hardware, and deployment conditions.
5	Jocher et al. / Ultralytics [6]	YOLOv8 implementation	Ultralytics software and model	Provides practical detection tools.
			family	Software citation, not a peer-reviewed accident-detection evaluation.
6	Bewley et al., 2016 [7]	Online multi-object tracking	SORT; Kalman filter and data association	Efficient tracking baseline; association can fail with occlusion or similar nearby objects.
7	Wojke et al., 2017 [8]	Tracking through occlusion	Deep SORT with appearance descriptors	Appearance cues can improve identity continuity; adds computation and depends on appearance quality.
8	Shashirangana et al., 2021 [9]	ALPR methods and challenges	Survey of localization, segmentation, and recognition	Highlights real-world ALPR challenges. Robustness remains difficult under blur, weather, and plate variation.
9	Khan et al., 2021 [10]	Automatic number-plate recognition	Survey and comparison of ANPR algorithms	Summarizes approaches for intelligent transportation; results across studies may not be directly comparable because datasets differ.
10	OASIS, 2019 [5]	IoT messaging standard	MQTT Version 5.0 specification	Defines lightweight publish/subscribe messaging; protocol alone does not validate incidents or guarantee emergency action.
11	Fang et al., 2024 [11]	Vision-based accident detection and anticipation survey	Review of vision-based methods, datasets, and evaluation metrics	Discusses rare and unpredictable accident events and dataset/evaluation challenges; it is a survey rather than a single deployed system.
12	Sahraei & Al Mamari, 2025 [12]	IoT accident detection and emergency notification review	Review of 101 peer-reviewed studies from 2008–2025	Compares sensors, integration frameworks, and recognition techniques; reports concerns including sensor accuracy, cost, weather, and standardization.

13	Razi & Chen, 2023 [13]	Deep learning for traffic-safety analysis	Review of traffic-video analysis pipelines and deep-learning tasks	Connects video processing, tracking, trajectory analysis, and anomaly detection; broad traffic-safety scope is not limited to accident notification.
14	Chai et al., 2024 [14]	Non-visual machine-learning approaches to road-accident prediction	Review of 95 studies on risk, occurrence, frequency, and severity	Shows the value of non-visual data for prediction; prediction tasks differ from real-time crash detection.
15	Prabu & Namitha, 2024 [15]	IoT-based automatic road accident detection review	Review of sensors, communication protocols, algorithms, and system architectures	Highlights integration and scalability issues; review findings depend on the included studies and implementation contexts.

2.3 Comparative Analysis

Vision-based studies provide scene context and can identify interactions that are difficult to capture with a single impact sensor. Temporal methods, such as anomaly modeling and trajectory-conflict analysis, are particularly relevant because accidents unfold over time [1], [2]. However, video analytics can be affected by low resolution, adverse weather, occlusion, camera placement, and unusual but non-accidental movements. A detector trained to identify vehicles cannot, by itself, prove that a collision has occurred.

Tracking algorithms can provide continuity across video frames, but tracking failures can propagate into trajectory analysis [7], [8]. ALPR adds a potentially useful identification channel, yet its literature reports persistent challenges from motion blur, lighting, viewing angle, and regional plate differences [9], [10]. Consequently, a system should attach confidence scores to detections and avoid treating an uncertain plate reading as definitive.

IoT sensors and GPS can provide independent evidence, while MQTT can transport structured incident messages [5]. The reviewed technical foundations support the feasibility of an integrated architecture, but they do not prove that combining all modules will always reduce false positives or response times. Those claims must be tested using representative accident and non-accident events, realistic communication conditions, and explicit evaluation criteria.

3. Applications of AI-Based Accident Detection and Emergency Response

3.1 Video-Based Vehicle and Event Detection

A camera stream can be sampled into frames and processed by a detector such as YOLOv8 to locate vehicles and other road users. Object positions are then passed to a tracking module, which estimates movement across frames. A downstream event classifier can evaluate abrupt trajectory changes, proximity, direction, and temporal patterns. Because a near miss or sudden stop may resemble an accident, this stage should output a probability or confidence level rather than an unqualified binary decision.

3.2 IoT Sensor Validation and GPS Localization

Accelerometers and gyroscopes mounted in a vehicle or connected to roadside infrastructure can identify sudden impacts or abnormal motion. A sensor event can trigger video review or provide a second evidence stream for the event classifier. GPS coordinates can then be associated with a timestamp and transmitted with the alert. Sensor thresholds must be calibrated because potholes, speed bumps, hard braking, and device movement may create false triggers.

3.3 Vehicle Identification Using ALPR and OCR

When image quality permits, a plate-localization model can identify the license plate region, after which OCR can read the characters. Plate recognition should be optional when the plate is blurred, occluded, distant, or poorly illuminated. The alert should retain the confidence score and, where legally appropriate, a cropped evidence image. Access to vehicle-identifying data should be restricted, logged, and governed by applicable privacy rules.

3.4 Communication and Edge–Cloud Processing

An edge device can perform initial detection and event verification near the camera or sensor, reducing the need to transmit all raw video continuously. After verification, the edge service can publish a compact incident message through MQTT to a broker or backend API. The cloud layer can store incident metadata, display a monitoring dashboard, and support later analysis. Reliable operation requires secure credentials, encrypted transport, timestamps, message identifiers, duplicate handling, acknowledgements, retry logic, and a clear procedure for failed delivery.

3.5 Severity Assessment and Emergency Notification

Severity classification may use available evidence such as the number of vehicles involved, estimated collision intensity, vehicle motion, scene context, and sensor measurements. Such classifications should initially be treated as decision-support estimates, not medical diagnoses. Emergency notifications should include the location, event time, confidence, available vehicle details, and a link to evidence when permitted. Integration with real emergency services requires authorization, operational agreements, and testing; a prototype should not send unverified live alerts to public emergency channels.

4. Benefits

An integrated AI–IoT approach may support continuous monitoring and reduce dependence on human observers watching video feeds. Combining visual context with physical sensor measurements can provide more evidence than either source alone. GPS can help locate an incident, while ALPR/OCR may contribute vehicle identification when the image is readable. Edge processing can reduce communication overhead and support faster local decisions, whereas cloud storage can enable centralized incident records, audit trails, and performance analysis.

The benefits are conditional rather than automatic. Fusion improves reliability only when the input signals are informative and the fusion logic handles missing or contradictory evidence. A false alert may consume emergency resources, while a missed event can delay assistance. Therefore, system benefits should be reported using measured detection quality, latency, message delivery reliability, and operational usability—not only by describing the proposed architecture.

5. Challenges

5.1 Environmental and Data Challenges

Traffic video varies by camera angle, resolution, weather, lighting, traffic density, and occlusion. Datasets may underrepresent severe accidents, rare vehicle types, night scenes, and unusual road layouts. Training and testing should be separated by video, location, or time where possible to reduce leakage between nearly identical frames. Evaluation should include normal traffic, sudden braking, road bumps, congestion, and near misses as negative or hard-negative examples.

5.2 False Positives and Severity Classification

A sudden change in vehicle trajectory does not necessarily mean a collision. Sensor thresholds can also be activated by non-accident impacts. Multimodal verification can help, but only if sources are sufficiently independent and synchronized. Severity classification is difficult because camera observations may not reveal occupant injuries or actual damage. The system should distinguish observed scene severity from confirmed human injury.

5.3 ALPR, Communication, and Infrastructure

ALPR performance may degrade under poor lighting, motion blur, unusual plate formats, and partial obstruction [9], [10]. Communication can fail because of poor connectivity, broker outages, or device power loss. Edge devices have limited processing and storage, while cloud services introduce cost, connectivity, and data-governance concerns. Message retries and local buffering can help, but duplicate and stale alerts must be handled safely.

5.4 Privacy and Security

Video footage and license plate numbers can be sensitive personal or identifying information. The system should use access control, encrypted communication, audit logging, retention limits, and purpose limitation. Only the minimum necessary data should be transmitted, and access to stored video should be restricted. Deployment must comply with applicable laws and institutional requirements, and emergency-service integration must be reviewed before operational use.

6. Research Gap

The reviewed sources suggest several areas for further investigation. These are framed as research opportunities to be validated against the literature rather than as claims that no prior system has addressed them. This is consistent with recent IoT accident-notification reviews, which highlight multi-sensor reliability and communication standardization as open implementation concerns [12], [15].

1. Limited integration of visual and IoT data: an integrated design must define how evidence is synchronized, weighted, and reconciled.
2. False positives in standalone systems: evaluation is needed on hard negatives such as sudden braking, road bumps, near misses, and dense traffic.
3. Insufficient severity classification: a practical system should separate collision confidence from estimated severity and avoid implying medical certainty.
4. Vehicle identification limitations: ALPR/OCR should provide confidence-aware identification and safe fallback behavior when the plate is unreadable.
5. Lightweight real-time communication: end-to-end latency, delivery acknowledgements, retries, and duplicate-alert handling require measurement.
6. Limited edge–cloud scalability: work allocation between edge and cloud should be evaluated for latency, bandwidth, cost, and resilience.
7. Incomplete emergency-response pipeline: detection must be connected to verification, location, evidence packaging, authorized notification, and acknowledgement.

The principal research opportunity is an end-to-end, confidence-aware architecture that follows: Detection → Verification → Identification → Severity Assessment → Location Tracking → Communication → Emergency Alert. Its contribution should be demonstrated through reproducible evaluation rather than assumed from component descriptions.

7. Proposed Framework and Methodology

7.1 Proposed System Architecture

The literature supports a modular architecture with the following stages:

Layer 1 — Data Acquisition: collect traffic video and, where available, accelerometer/gyroscope readings and GPS coordinates. Each record should include a timestamp and source identifier.

Layer 2 — Computer Vision: use an object detector such as YOLOv8 to detect road users. A tracker such as SORT or Deep SORT associates detections across frames and produces movement trajectories [6]–[8].

Layer 3 — Accident Candidate Detection: evaluate temporal changes, trajectory conflicts, object proximity, abrupt motion, and visual evidence. The output is a candidate event with a confidence score, not an unquestioned confirmation.

Layer 4 — Sensor-Based Verification: compare the video-based candidate with available sensor evidence. Missing or contradictory signals should be recorded rather than silently treated as confirmation.

Layer 5 — Vehicle Identification and Location: use ALPR/OCR only when image quality is adequate; attach recognition confidence. Associate GPS coordinates and event time where available.

Layer 6 — Alert Communication: send a structured event message using MQTT or a secure backend API. Include a unique incident identifier, timestamp, coordinates, confidence, status, and permitted evidence references.

Layer 7 — Dashboard, Audit, and Human Review: store event metadata, display alerts to an authorized reviewer, log decisions, and track delivery acknowledgements. Severe or uncertain cases should follow a defined human-review procedure.

7.2 Suggested Incident Message

A structured message may contain: incident_id, event_timestamp, latitude, longitude, detection_confidence, severity_estimate, vehicle_plate_text (optional), plate_confidence (optional), evidence_reference, delivery_status, and review_status. Sensitive fields should be omitted when unnecessary and protected through access control and retention rules.

7.3 Evaluation Methodology

A future prototype should be tested on video clips separated by scene or recording session to reduce data leakage. The test set should include accidents, near misses, abrupt braking, congestion, occlusion, night-time footage, and ordinary traffic. Sensor tests should include both accident-like impacts and non-accident disturbances. Communication tests should include intermittent connectivity, duplicate messages, broker interruption, and delayed acknowledgements.

Recommended metrics include precision, recall, F1-score, false alerts per hour, missed-event rate, time from event onset to detection, end-to-end notification latency, ALPR character-level accuracy, message-delivery success rate, and resource usage on the target device. Results should be reported for individual modules and the complete pipeline. No experimental scores are reported in this literature review because no project test results were supplied.

8. Future Scope

Future work can implement and evaluate a prototype using a suitable traffic-video dataset and edge-capable hardware. A semantic temporal classifier could be compared with rule-based trajectory analysis. Sensor fusion could be tested under missing, delayed, and contradictory inputs. ALPR should be evaluated separately across day, night, blur, viewing angle, and plate-region variation. The communication layer could be assessed for latency, delivery reliability, reconnection, and duplicate prevention. Additional work may explore privacy-preserving video storage, edge inference optimization, explainable incident evidence, and authorized integration with emergency-response organizations. A public or institutionally approved benchmark with clearly defined accident and non-accident labels would improve reproducibility.

9. Conclusion

This literature review examined technical foundations for AI-based road accident detection and emergency response using computer vision and IoT. Video anomaly detection and trajectory-conflict analysis show how temporal behavior can support accident-candidate detection [1], [2]. Object detection and tracking provide useful components, but they do not independently prove that a collision has occurred [3], [4], [7], [8]. ALPR may contribute vehicle identification, although image conditions and plate variation can reduce reliability [9], [10]. MQTT provides a lightweight communication mechanism, but dependable emergency notification also requires application-level security, acknowledgement, retries, duplicate handling, and operational authorization [5].

The review therefore motivates a modular pipeline combining video analytics, optional sensor evidence, confidence-aware vehicle identification, GPS localization, secure messaging, and human oversight. Its practical value remains to be established through reproducible tests of detection quality, false alarms, latency, communication reliability, privacy, and usability. The proposed architecture should be treated as a design direction for future implementation, not as an experimentally validated system.

References

[1] K. Pawar and V. Attar, “Deep learning based detection and localization of road accidents from traffic surveillance videos,” ICT Express, vol. 8, no. 3, pp. 379–387, 2022. doi: 10.1016/j.icte.2021.11.004.

[2] H. Ghahremannezhad, H. Shi, and C. Liu, “Real-Time Accident Detection in Traffic Surveillance Using Deep Learning,” in Proc. IEEE International Conference on Imaging Systems and Techniques (IST), 2022. doi: 10.1109/IST55454.2022.9827736.

[3] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, “You Only Look Once: Unified, Real-Time Object Detection,” in Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016, pp. 779–788. doi: 10.1109/CVPR.2016.91.

[4] A. Bochkovskiy, C.-Y. Wang, and H.-Y. M. Liao, “YOLOv4: Optimal Speed and Accuracy of Object Detection,” arXiv:2004.10934, 2020.

[5] OASIS, MQTT Version 5.0, OASIS Standard, Mar. 2019. Available: https://docs.oasis-open.org/mqtt/mqtt/v5.0/mqtt-v5.0.html

[6] G. Jocher et al., “Ultralytics YOLO,” software documentation and repository. Available: https://github.com/ultralytics/ultralytics. Accessed Oct. 12, 2026.

[7] A. Bewley, Z. Ge, L. Ott, F. Ramos, and B. Upcroft, “Simple Online and Realtime Tracking,” in Proc. IEEE International Conference on Image Processing (ICIP), 2016, pp. 3464–3468. doi: 10.1109/ICIP.2016.7533003.

[8] N. Wojke, A. Bewley, and D. Paulus, “Simple Online and Realtime Tracking with a Deep Association Metric,” in Proc.

IEEE International Conference on Image Processing (ICIP), 2017, pp. 3645–3649. doi: 10.1109/ICIP.2017.8296962.

[9] J. Shashirangana, H. Padmasiri, D. Meedeniya, and C. Perera, "Automated License Plate Recognition: A Survey on Methods and Techniques," IEEE Access, vol. 9, pp. 11203–11225, 2021. doi: 10.1109/ACCESS.2020.3047929.

[10] M. A. Khan et al., "Automatic Number Plate Recognition: A Detailed Survey of Relevant Algorithms," Sensors, vol. 21, no. 9, article 3028, 2021. doi: 10.3390/s21093028.

[11] J. Fang, Y. Qiao, J. Xue, and Z. Li, "Vision-Based Traffic Accident Detection and Anticipation: A Survey," IEEE Transactions on Circuits and Systems for Video Technology, vol. 34, no. 4, pp. 1983–1999, Apr. 2024, doi: 10.1109/TCSVT.2023.3307655.

[12] M. A. Sahraei and S. R. M. Al Mamari, "A Review of Internet of Things Approaches for Vehicle Accident Detection and Emergency Notification," Sustainability, vol. 17, no. 14, Art. no. 6510, 2025, doi: 10.3390/su17146510.

[13] A. Razi and X. Chen, "Deep learning serves traffic safety analysis: A forward-looking review," IET Intelligent Transport Systems, vol. 17, no. 1, pp. 22–71, 2023, doi: 10.1049/itr2.12257.

[14] A. B. Z. Chai, B. T. Lau, M. K. T. Tee, and C. McCarthy, "Enhancing road safety with machine learning: Current advances and future directions in accident prediction using non-visual data," Engineering Applications of Artificial Intelligence, vol. 137, Art. no. 109086, 2024, doi: 10.1016/j.engappai.2024.109086.

[15] C. R. Prabu and K. Namitha, "IoT-based automatic road accident detection systems: A review," in Proc. 4th International Conference on Soft Computing for Security Applications (ICSCSA), 2024, doi: 10.1109/ICSCSA64454.2024.00122