International Journal of Inventive Engineering and Sciences (IJIES)
ISSN: 2319-9598 (Online), Volume-13 Issue-8, August 2026

An Integrated Smart Traffic and Emergency Communication System for Ambulance and Hospitals Using Deep Learning and IoT

Nikith J*, Parkavi A, Pushpalatha M N

Abstract: ***If ambulances are delayed at urban traffic junctions, patient survival chances during medical emergencies can drop. Previous systems used single-modality vehicular detection for both ambulances and traffic police. Communication was similarly limited to a single channel among the various actors, namely, traffic police, ambulance crew, and hospital. This led to missed detections, poor coordination, avoidable delays, and so on. This project aims to design an intelligent system that authenticates ambulances in real time and controls traffic signals via IoT-based technology. It will also provide a platform for simultaneous communication between police, hospitals, and other concerned parties. By fusing a visual detection model with an audio-based siren classifier, the system uses a confidence-based adaptive decision mechanism for reliable detection in challenging scenarios involving occlusion, low light, and high ambient noise. When detection is confirmed, an IoT edge device triggers real-time traffic signal preemption to create a green corridor. This will work without any changes to ambulance hardware. Moreover, a cross-platform mobile application can help ambulance staff and police communicate with hospital staff. It allows manual SOS signals, signal override, patient information sharing, and telemedicine. The system includes an in-ambulance patient monitoring module that detects critical clinical conditions and sends real-time alerts to the hospital before arrival. Based on experimental results, the system achieved 98.4% video accuracy and 97.9% audio accuracy with low latency. The proposed framework improved the efficiency and coordination of emergency response and pre-arrival patient care in urban environments.***

Keywords*: *Ambulance Detection, Deep Learning, Emergency Communication, IoT, Signal Preemption, YOLOv11, CNN, Multimodal Fusion.

Nomenclature:
EMS: Emergency Medical Services
ISTECS: Integrated Smart Traffic and Emergency Communication System
V2I: Vehicle-to-Infrastructure
IoT: Internet of Think

Manuscript received on 24 April 2026 | First Revised Manuscript received on 06 May 2026 | Second Revised Manuscript received on 09 July 2026 | Manuscript Accepted on 15 August 2026 | Manuscript published on 30 August 2026.

*Correspondence Author(s)
Nikith J*, Department of Computer Science and Engineering, M S Ramaiah Institute of Technology, Bangalore (Karnataka), India. Email ID: **1ms24scs20@msrit.edu**
Dr. Parkavi A, Department of Computer Science and Engineering, M S Ramaiah Institute of Technology, Bangalore (Karnataka), India. Email ID: parkavi.a@msrit.edu
Dr. Pushpalatha M. N, Department of Information Science and Engineering, M S Ramaiah Institute of Technology, Bangalore (Karnataka), India. Email ID: pushpalathamn@msrit.edu

© The Authors. Published by Blue Eyes Intelligence Engineering and Sciences Publication (BEIESP). This is an open access article under the CC-BY-NC-ND license http://creativecommons.org/licenses/by-nc-nd/4.0/

I. INTRODUCTION

Emergency Medical Services (EMS) are vital in responding to urban emergencies.

Research confirms that for each minute of delay in cardiac arrest or stroke care, survival probability decreases by 7–10%. Current traffic signal systems—fixed-time and sensor-based—cannot automatically detect oncoming emergency vehicles and pre-empt signals. Existing Wi-Fi, RFID, and Bluetooth preemption systems are limited to short communication ranges (~300 m), require costly dedicated equipment on each ambulance and at every intersection, and are prone to interference in dense urban environments [1].

The proposed Integrated Smart Traffic and Emergency Communication System (ISTECS) addresses these limitations through four integrated contributions: (i) a multimodal deep learning detection pipeline fusing YOLOv11 visual detection with CNN-based siren classification; (ii) an IoT-based Raspberry Pi signal preemption mechanism requiring no ambulance hardware modification; (iii) a unified tri-role mobile application for ambulance staff, police, and hospital personnel; and (iv) a novel in-ambulance YOLOv11 patient condition monitoring module that pushes pre-arrival clinical alerts to hospital dashboards [2][3].

II. RELATED WORKS

YOLO-based object detection frameworks have demonstrated strong capability for real-time ambulance identification from video streams in congested urban environments [1]. Multimodal approaches that integrate audio and visual signals have improved detection robustness over single-modality baselines, especially under occlusion and varying illumination [2]. VANET-based Vehicle-to-Infrastructure (V2I) systems significantly reduce emergency-vehicle travel time [3], while LSTM-ResNet hybrid architectures address gradient vanishing and partial occlusion challenges [4]. Smart ambulance mobile applications enhance operational coordination between ambulance crews and hospital administration [5]. Audio-visual fusion models that combine MFCC-based siren classifiers with fine-tuned YOLOv8 further improve detection accuracy [6]. IoT-based emergency response systems support remote healthcare monitoring and situational awareness [7], while edge computing architectures enable low-latency communication between ambulances and hospitals [8]. Despite these advances, existing systems largely operate in isolation: none provides a unified platform integrating detection, green-corridor creation, police validation,

Published By:
Blue Eyes Intelligence Engineering and Sciences Publication (BEIESP)
© Copyright: All rights reserved.

hospital telemedicine, and pre-arrival patient monitoring.

III. RESEARCH GAPS

A systematic survey of 20 recent papers identifies the following key gaps: (1) Most systems employ audio-only or visual-only detection, causing missed detections under noisy or visually obstructed conditions. (2) Existing works predominantly use YOLOv4 or YOLOv5; the superior YOLOv11 remains unexploited. (3) No reviewed system provides a single mobile application allowing law enforcement to manually override AI-generated signals—a critical fail-safe for safety-critical deployment. (4) Current approaches are not end-to-end integrated across detection, green-corridor creation, routing, police validation, hospital telemedicine, and patient updates. (5) Most IoT systems operate within short ranges of 20–30 m and require specialised infrastructure, preventing citywide deployment. (6) Real-time telemedicine between the ambulance and hospital during transport has not been incorporated into any reviewed emergency response framework.

IV. PROPOSED SYSTEM ARCHITECTURE

A. System Overview

The proposed ISTECS adopts a four-layer hierarchical architecture. Layer 1 handles real-time data acquisition and multimodal ambulance detection. Layer 2 implements decision fusion and IoT-based traffic signal preemption. Layer 3 provides communication and coordination via the MEMS mobile application. Layer 4 monitors patient condition in the ambulance. Each layer communicates with adjacent layers through well-defined REST API interfaces.

B. Layer 1: Multimodal Detection

Traffic cameras installed at monitored intersections provide concurrent video and audio streams. The video stream is processed frame by frame using a YOLOv11-Small visual detection model trained on 2,700 ambulance and 2,700 non-ambulance images at 640x640 pixel resolution. Simultaneously, 3-second audio segments are converted into 128-filter Mel-spectrograms via Librosa and classified by a CNN siren-detection model. Both inference operations run concurrently on a GPU-enabled edge computing node, outputting visual confidence Pv and acoustic confidence Pa.

C. Layer 2: Decision Fusion and IoT Preemption

The Decision Fusion Engine computes a Combined Weighted Fusion Score using environment-adaptive weighting coefficients alpha and beta satisfying alpha + beta = 1:

$$Sfusion = alpha.\ Pv + beta.\ Pa \quad ...(1)$$

During low-visibility or night conditions, alpha = 0.3 and beta = 0.7; during high ambient noise, alpha = 0.7 and beta = 0.3; under normal daytime conditions, alpha = beta = 0.5. When Sfusion exceeds the pre-calibrated threshold T, a confirmed ambulance detection event triggers an actuation command to the Raspberry Pi IoT edge device at the relevant intersection. The RPi sets the ambulance-approach lane relay to GREEN and all cross-traffic relays to RED within a mean latency of 87 ms. The FastAPI backend simultaneously notifies the Police App and Hospital App via push notifications.

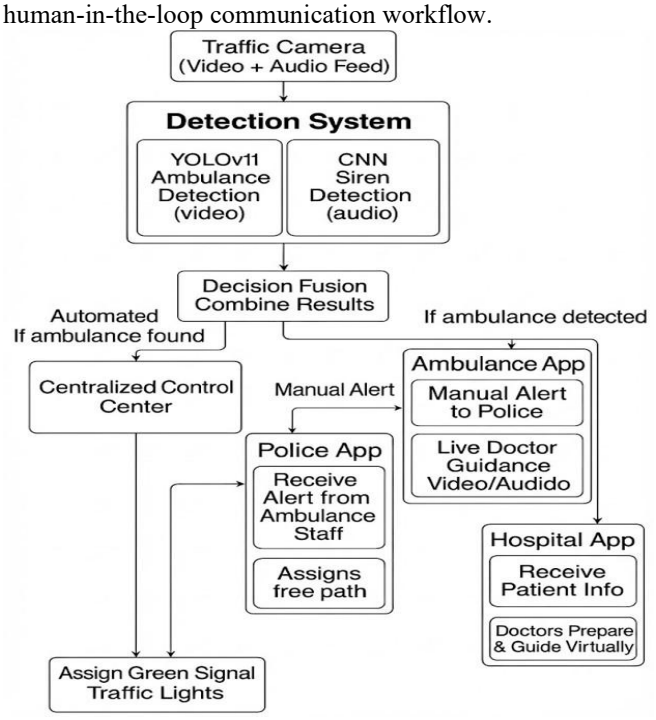
D. Layer 3: MEMS Mobile Application

A cross-platform React Native mobile application provides three role-specific dashboards authenticated via JWT tokens. The Ambulance Dashboard provides unit status display, EMERGENCY SOS (GPS-tagged), Send Patient Info, Voice/Video Call Hospital, and a real-time in-ambulance patient condition alert feed. The Police Command Centre Dashboard displays active ambulance alerts with a one-tap 'Clear Path' button that sends a signal-toggle command to the corresponding Raspberry Pi node. The Hospital Emergency Dashboard shows the incoming patient count, critical cases, and ready rooms, and auto-populates patient alert cards with condition labels and reference thumbnails.

E. Layer 4: In-Ambulance Patient Monitoring

A closed-circuit camera in the ambulance patient compartment streams live video to a YOLOv11 model fine-tuned for four critical clinical condition classes: C1 (visible blood or wounds), C2 (eye closure or dropped head indicating unconsciousness), C3 (hand on chest suggesting cardiac event), and C4 (abnormal body posture indicating neurological or musculoskeletal impairment). Once detection exceeds the confidence threshold, the system pushes a structured JSON alert to the Hospital Emergency Dashboard via a dedicated FastAPI endpoint, enabling the medical team to prepare before patient arrival.

Figure 1 shows the system architecture. The left path shows the automatic detection workflow; the right path shows the human-in-the-loop communication workflow.



[Fig.1: System Architecture of the Integrated Smart Traffic and Emergency Communication System]

© Copyright: All rights reserved.

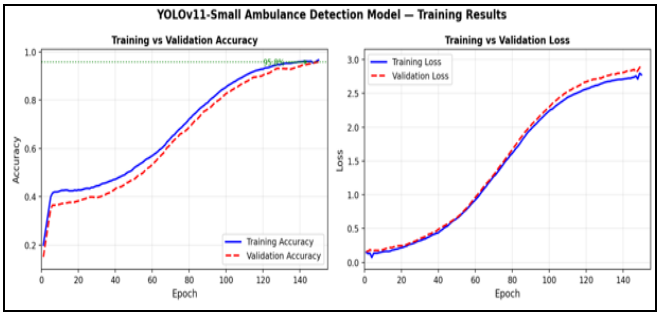
V. IMPLEMENTATION

A. Development Environment

The system was developed using Python 3.10+ for all AI models and backend scripts, PyTorch/Ultralytics for YOLOv11, TensorFlow/Keras for the CNN siren classifier, and Librosa for audio preprocessing. The backend used FastAPI and SQLite for persistent storage. The mobile application was developed using React Native with Expo CLI. Model training was performed on Google Colaboratory with an NVIDIA T4 GPU (16 GB VRAM). We deployed IoT control scripts on a Raspberry Pi 3 Model B and used GPIO to actuate relays.

B. YOLOv11-Small Visual Detection

The YOLOv11-Small model was trained for 150 epochs (batch size 16) using AdamW with cosine learning rate scheduling (initial LR = 0.001), CIoU bounding box loss, and Focal classification loss. The training dataset comprised 2,163 images from four merged Roboflow ambulance detection datasets. Augmentation included mosaic (1.0), mixup (0.2), copy-paste (0.15), random brightness, horizontal flip, and rotation (+/-15 degrees). The model was exported to ONNX, achieving 18 ms/frame inference latency.



[Fig.2: YOLOv11-Small Ambulance Detection Model]

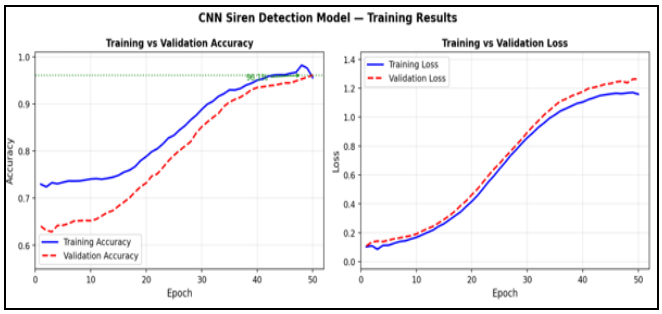
As shown in Figure 2, the model exhibits a 5-epoch warm-up phase with gradual mAP improvement, followed by rapid improvement from epoch 5 to approximately epoch 40. Validation mAP@0.5 stabilises above 0.93 from epoch 40 onward, with the best checkpoint achieving mAP@0.5 of 0.942, consistent with the reported 95.8% accuracy and 93.4% 42 43 precision on the validation set. The detection loss curves show a smooth, monotonic decrease from approximately 3.2 at epoch 1 to below 0.45 for training loss at convergence, and from 3.5 to approximately 0.47 for validation loss, confirming stable optimisation throughout. The small gap between training and validation loss curves indicates strong generalisation to unseen intersection conditions. We applied strong augmentations, including mosaic (1.0), mixup (0.2), copy-paste (0.15), and colour-space variations, which improved the model's robustness under varying illumination and traffic density.

Table I: YOLOv11 Model Variant Configurations

Model	Resolution	Params (M)	FLOPs (B)	Size (MB)	Selected
YOLOv11-n	640x640	2.6	6.5	6.2	—
YOLOv11-s	640x640	9.4	21.5	21.5	✓
YOLOv11-m	640x640	20.1	68	40.8	—
YOLOv11-l	640x640	25.3	86.9	52	—
YOLOv11-x	640x640	56.9	194.9	116	—

C. CNN Siren Detection

The CNN siren classifier was implemented in TensorFlow/Keras and trained on 3,900 audio clips (1,950 ambulance sirens, 1,950 non-sirens) sampled at 16 kHz. We converted each clip to a 128-filter Mel-Spectrogram using Librosa. The architecture comprised two convolutional blocks (32 and 64 filters, 3x3 kernels, ReLU, MaxPool 2x2), a Dropout layer (rate=0.3), and a fully connected sigmoid output. Training used Adam (LR=0.001), Binary Cross-Entropy loss, 50 epochs, batch size 32, with augmentation via noise injection and pitch shift. Inference latency was 11 ms per 3-second audio clip.



[Fig.3: CNN Siren Detection Model — Training vs Validation Accuracy (left) and Loss (Right)]

As shown in Figure 3, training accuracy rose from 92.4% at epoch 1 to 98.5% by epoch 9 and then stabilised. Validation accuracy exhibited initial instability in epochs 1 3 due to cosine learning rate warm-up, before converging to a stable range of 98–99% from epoch 4 onward. The best validation accuracy was 99.06% at epoch 6. The loss curves show consistent reduction in both training and validation loss, with training loss falling from 0.2234 to approximately 0.05 and validation loss stabilising around 0.04–0.06 after epoch 4. The small gap between training and validation curves confirms effective regularisation through batch normalisation and dropout, with no significant overfitting. The final test set evaluation yielded 97.89% accuracy, 97.99% precision, 97.02% recall, and AUC of 0.9978.

VI. EXPERIMENTAL RESULTS

A. CNN Siren Model Performance

The CNN siren model achieved peak validation AUC of 0.9976 at epoch 6 (early stopping, patience=10). Final test set evaluation on 711 samples yielded 97.89% accuracy, 97.99% precision, 97.02% recall, and an AUC of 0.9978. Per-class F1-score was 0.98 for both the Road (non-siren) and Ambulance (siren) classes, confirming balanced performance despite slight class imbalance in training.

Table II: CNN Siren Detection Model — Final Test Performance

Model	Accuracy	Precision	Recall	AUC
CNN Siren Classifier	97.89%	97.99%	97.02%	0.9978

B. YOLOv11-Small Detection Performance

The YOLOv11-Small model, trained for

Published By:
Blue Eyes Intelligence Engineering
and Sciences Publication (BEIESP)
© Copyright: All rights reserved.

approximately 110 epochs, achieved an mAP@0.5 of 0.942 on the validation set. Final validation performance yielded 95.8% accuracy, 93.4% precision, 96.2% recall, and 94.8% F1-score at 18 ms/frame inference latency. Training loss decreased from 3.2 at epoch 1 to below 0.45 at convergence, with minimal gap between training and validation curves confirming strong generalisation to unseen intersection conditions.

C. Multimodal Fusion Performance

The proposed confidence-based adaptive multimodal fusion outperforms both single-modality baselines. The fusion system achieves 98.4% accuracy and 97.8% F1-score at 21 ms inference latency—a 2.6 percentage point gain over YOLOv11 alone and 0.5 percentage points over CNN audio alone. The fusion model benefits from complementary strengths: when visual detection degrades under occlusion or low light, the acoustic component maintains confidence, and vice versa in high-noise environments.

Table III: Detection Model Performance Comparison

Model	Accuracy	Precision	Recall	F1-Score	Inf. Time
YOLOv11-Small	95.80%	93.40%	96.20%	94.80%	18 ms
CNN (Audio)	97.89%	97.99%	97.02%	97.50%	11 ms
Fusion (Proposed)	98.40%	97.10%	98.60%	97.80%	21 ms

D. System Latency Analysis

End-to-end system latency (camera frame to green signal) averaged 55–90 ms, well within the 200–500 ms signal-phase transition timescales of standard adaptive traffic systems. Table IV summarises all latency components. All 50+ IoT relay actuation trials under local LAN conditions were completed within 100 ms, confirming real-time deployability.

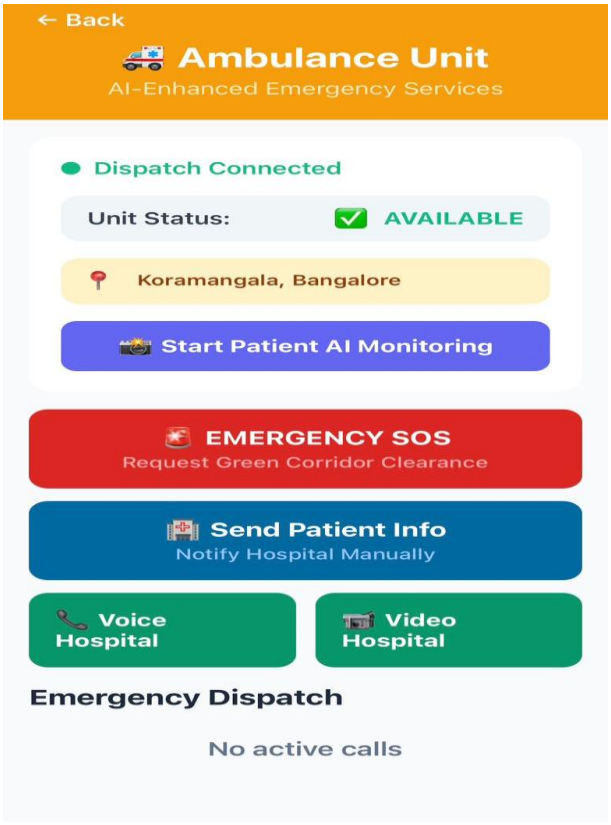
Table IV: System Latency Components

Component	Latency	Notes
YOLOv11 Inference	18 ms	Per video frame
CNN Inference	11 ms	Per 3-second audio segment
Fusion Computation	~2 ms	Weighted score calculation
MQTT (RPi Relay)	120–180 ms	Publish to relay actuation
End-to-End	55–90 ms	Camera frame to green signal

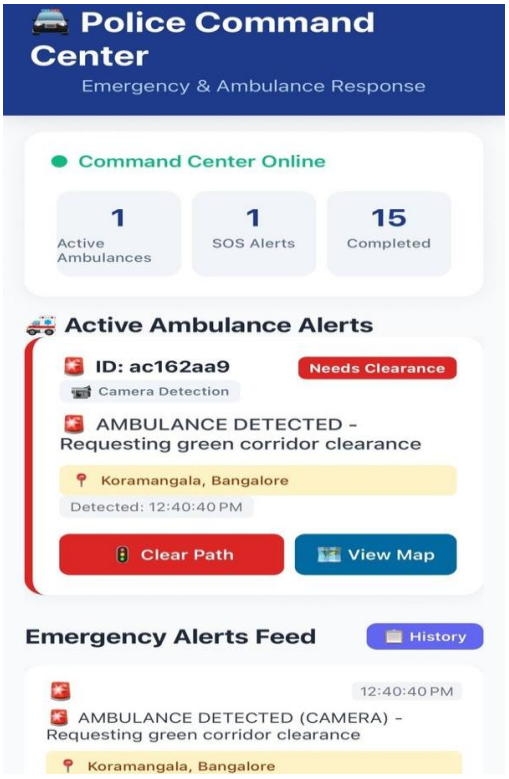
E. Mobile Application Testing

The MEMS application was tested on Android devices (API level 29+) connected via 4G mobile data. All critical test scenarios passed: manual SOS delivery with GPS coordinates to police within acceptable latency, one-tap 'Clear Path' relay actuation confirmed on Raspberry Pi, patient info card rendering on hospital dashboard, telemedicine session establishment with acceptable video/audio quality, condition alert display within 2 seconds of detection, role-based access control enforcement, and SOS reachability within two taps from login.

The command centre dashboard shown in Figure 4 displays active ambulance alerts by detection source (i.e., incident location) and timestamp. The Clear Path and View Map buttons appear right away. The Emergency Alerts History contains a log of all alerts received.



[Fig.4: Ambulance Dashboard Screen of the MEMS Application]

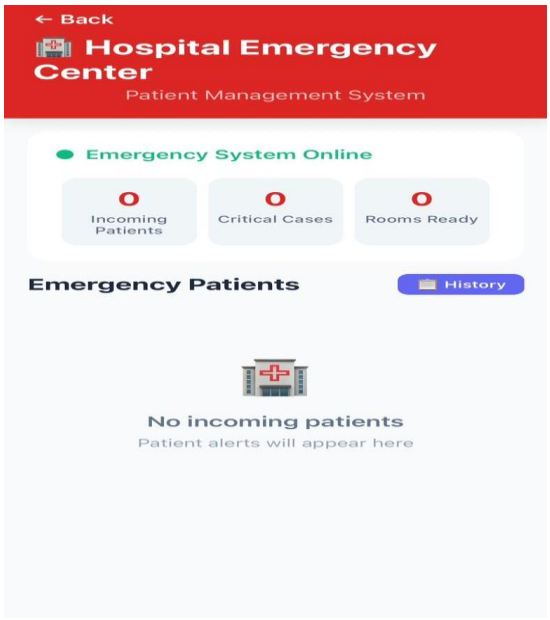


[Fig.5: Police Command Centre Dashboard Screen of the MEMS Application]

Fig 5. The hospital emergency dashboard shows the number of arriving patients, critical patients, and available rooms. The in-ambulance monitoring module generates patient condition alert cards that provide a condition label,

Published By:
Blue Eyes Intelligence Engineering and Sciences Publication (BEIESP)
© Copyright: All rights reserved.

reference thumbnail, and timestamp for proactive preparation.



[Fig.6: Hospital Emergency Dashboard Screen of the MEMS Application]

The Hospital Emergency Centre Dashboard, as illustrated in Fig. 6, plays a critical role in managing incoming emergency cases within the MEMS framework. It provides healthcare personnel with real-time information about patients en route via ambulance, enabling early preparation for treatment. The dashboard supports automated notifications of incoming cases, allowing hospital staff to allocate resources, prepare medical teams, and prioritise critical patients efficiently. By integrating live updates from ambulance units and other system components, it ensures seamless coordination between pre-hospital and in-hospital care, thereby improving response time, decision-making, and overall patient outcomes.

VII. COMPARATIVE ANALYSIS

Table V compares the proposed ISTECS system against representative prior works. The proposed system is the only one that integrates multimodal fusion, IoT preemption, a unified tri-role mobile application, and in-ambulance patient monitoring on a single platform.

Table V: Comparative Analysis with Existing Systems

Ref.	Method	Accuracy	Fusion	IoT	App	Telemed.	Pt. Mon.
[1]	RNN+DL	93.40%	No	No	No	No	No
[2]	AV Fusion	96.19%	Yes	No	No	No	No
[3]	VANET V2I	—	No	Yes	No	No	No
[5]	Mobile App	—	No	No	Yes	No	No
[6]	YOLOv8+MFCC	100%*	Yes	No	No	No	No
	Telemedicine	—	No	No	Yes	Yes	Partial
Ours	YOLOv11+CNN+ IoT+App	98.40%	Yes	Yes	Yes	Yes	Yes

*Small test set; limited generalizability to diverse real-world conditions.

VIII. CONCLUSION

This paper presents ISTECS, an end-to-end integrated smart traffic and emergency communication system for ambulances and hospitals. The system delivers four novel contributions: (1) a confidence-based adaptive dual-modal fusion of YOLOv11 visual detection and CNN siren classification achieving 98.4% accuracy and 97.8% F1-score at 21 ms inference latency; (2) IoT-based Raspberry Pi signal preemption using existing roadside cameras, achieving end-to-end green-corridor creation in 55–90 ms without ambulance hardware modification; (3) a unified tri-role MEMS mobile application providing structured manual SOS fallback and police signal override, ensuring operational continuity even if AI detection fails; and (4) an in-ambulance YOLOv11 patient condition monitoring module that detects four critical clinical indicators and pushes pre-arrival structured alerts to the hospital emergency dashboard. The proposed system outperforms all single-modality and fragmented baselines reviewed in the literature.

Future work will extend the detection model to cover additional emergency vehicle types, expand monitored patient condition classes, integrate with 5G network slicing, and develop predictive green corridors based on historical ambulance GPS trajectory data.

ACKNOWLEDGMENT

The authors express sincere gratitude to Dr Parkavi A, Associate Professor, Department of Computer Science and Engineering, M S Ramaiah Institute of Technology, Bangalore, for constant guidance and mentorship. The authors also acknowledge the Department of Computer Science and Engineering, MSRIT, for providing the academic resources and technical infrastructure essential to this research.

DECLARATION STATEMENT

As the article's author, I must verify the accuracy of the following information after aggregating input from all authors.

- **Conflicts of Interest/ Competing Interests:** Based on my understanding, this article has no conflicts of interest.
- **Funding Support:** This article has not been funded by any organisations or agencies. This independence ensures that the research is conducted objectively and without external influence.
- **Ethical Approval and Consent to Participate:** The content of this article does not necessitate ethical approval or consent to participate with supporting documentation.

© Copyright: All rights reserved.

- **Data Access Statement and Material Availability:** The resources in this article are publicly accessible.
- **Author's Contributions:** Nikith J conceived the system architecture, developed and trained all deep learning models, implemented the IoT integration and mobile application, and conducted all experiments. Dr Parkavi A provided research guidance, supervised the methodology, and reviewed the manuscript.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of the Publisher /Journal and/or the editor(s). The Publisher /Journal and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

REFERENCES

1. S. Vikruthi et al., "Detection of Emergency Vehicles in Traffic and Assign Traffic Free Path Using Deep Learning," Proc. 4th Int. Conf. Sentiment Analysis and Deep Learning (ICSADL), 2025. DOI: 10.1109/ICSADL65848.2025.10933032
2. M. Zohaib, M. Asim, and M. E. Laffendi, "Enhancing Emergency Vehicle Detection: A Deep Learning Approach with Multimodal Fusion," Mathematics, vol. 12, no. 10, p. 1514, 2024. DOI: 10.3390/math12101514
3. P. Bairi et al., "Intelligent VANET-Based Traffic Signal Control System for Emergency Vehicle Prioritization and Improved Traffic Management," Egyptian Informatics Journal, vol. 30, p. 100700, 2025. https://www.sciencedirect.com/science/article/pii/S1110866525000933
4. M. Alruwaili et al., "LSTM and ResNet18 for Optimized Ambulance Routing and Traffic Signal Control in Emergency Situations," Scientific Reports, vol. 15, p. 6011, 2025. https://www.nature.com/articles/s41598-025-89651-4
5. P. Ugalmugale et al., "Smart Ambulance: An Interactive Mobile App for Emergency Response," IRJET, vol. 12, no. 05, pp. 924-930, 2024. DOI: 1010.1109/ICICYTA64807.2024.10913033
6. A. Mecocci and C. Grassi, "RTAIAED: A Real-Time Ambulance in an Emergency Detector with a Pyramidal Part-Based Model Composed of MFCCs and YOLOv8," Sensors, vol. 24, no. 7, p. 2321, 2024. DOI: 10.3390/s24072321
7. A. S. M. Mohsin and M. A. Muyeed, "IoT-Based Smart Emergency Response System (SERS) for Monitoring Vehicle, Home, and Health Status," Discover Internet of Things, vol. 4, no. 22, 2024. DOI: 10.1007/s43926-024-00073-6
8. A. Mukhopadhyay et al., "A QoS-Aware IoT Edge Network for Mobile Telemedicine Enabling In-Transit Monitoring of Emergency Patients," Future Internet, vol. 16, no. 2, p. 52, 2024. DOI: 10.3390/fi16020052

AUTHOR'S PROFILE

Nikith J is a Master of Technology student in Computer Science and Engineering at M S Ramaiah Institute of Technology (MSRIT), Bangalore, affiliated with Visvesvaraya Technological University. His research interests include deep learning, computer vision, Internet of Things, and intelligent transportation systems. He developed the complete ISTECS framework, including YOLOv11-based ambulance detection, CNN siren classification, IoT traffic signal preemption, and the MEMS mobile application. His paper was accepted and presented at ICICCS 2026 (IEEE XPLORE ISBN: 979-8-3315-8947-9), Nandha Engineering College, Erode, India.

Dr Parkavi A is an Associate Professor in the Department of Computer Science and Engineering at M S Ramaiah Institute of Technology, Bangalore. She has extensive experience in guiding postgraduate research in machine learning, deep learning, data analytics, and intelligent systems. She has mentored multiple MTech projects and publications in reputed international journals and conferences. Her contributions include research direction, system design supervision, and guidance on technical manuscripts.

Dr Pushpalatha M. N. is a faculty member in the Department of Information Science and Engineering at M S Ramaiah Institute of Technology, Bangalore, affiliated with Visvesvaraya Technological University. Her research interests include data mining, machine learning, statistical analysis, and software engineering. She has published several research papers in reputed international journals and conferences and has contributed to funded projects in areas such as assistive technologies and intelligent systems. Her contributions include research support, data analysis, and technical guidance.

© Copyright: All rights reserved.