International Journal for Multidisciplinary Research (IJFMR)

E-ISSN: 2582-2160 ● Website: www.ijfmr.com ● Email: editor@ijfmr.com

Designing a Smart Emergency Framework for Indian Cities: A Data-Driven GIS Approach to Optimized Vehicle Dispatch and Routing

Sakshi Bhardwaj

Abstract

This study proposes a smart emergency response framework tailored for Indian cities, aiming to reduce ambulance delays through dynamic routing and real-time traffic management. Leveraging Geographic Information Systems (GIS), predictive analytics, and tools like OpenStreetMap and ISRO's Bhuvan, the framework integrates accident hotspot mapping, live congestion data, and green corridor activation. A Bangalore-based simulation showed that response times could be cut by up to 35%. The system architecture includes real-time GPS tracking, adaptive signal control, and a predictive dashboard. Results demonstrate the potential for scalable, data-driven reforms in emergency dispatch—shifting from reactive to proactive models that can save lives.

Chapter 1: Introduction

In Indian cities, the sound of a siren is meant to signal urgency—but too often, it fails to translate into swift action. Emergency vehicles like ambulances, fire trucks, and police vans frequently get stuck in the very problem they're meant to cut through: dense traffic, narrow lanes, and unplanned urban sprawl. With India's urban population growing rapidly, these delays aren't just frustrating—they can be deadly. A 2020 report found that the lack of real-time traffic coordination and basic infrastructure plays a major role in slowing down emergency response times (Moazum Wani, Khan, & Alam, 2020; Indian Awaaz, 2024).

Cities like Delhi and Mumbai are among the most congested in the world. Ambulances in these metros often take anywhere from 20 to 45 minutes to reach patients, well past the critical Golden Hour for trauma and cardiac emergencies (GoAid, 2024; PubMed, 2024). Sirens often go ignored. Illegal parking, reckless driving, and poorly designed roads all contribute to the problem, creating choke points that can turn life-saving missions into logistical nightmares.

Technology offers a potential way out, but its adoption is still uneven. Around the world, smart cities are using real-time sensors, geospatial intelligence, and dynamic routing systems to speed up emergency response. India has the tools—initiatives like the Smart Cities Mission and ISRO's Bhuvan platform lay the groundwork—but implementation lags behind. GIS (Geographic Information Systems) have yet to be fully integrated into live dispatch systems, limiting their impact on the ground (Rajkot pilot, Times of India, 2024; Moazum Wani et al., 2020).

Traditional dispatch models in India are still largely static, based on fixed zones and manual coordination. However, pilot projects are showing what's possible. In Rajkot, a GPS-linked traffic signal system allowed ambulances to cut through intersections by triggering green corridors automatically—cutting down on response times significantly (Times of India, 2024). What's needed now is a fully data-driven dispatch network—one that tracks real-time vehicle locations, traffic flow, and accident hotspots to suggest optimal routes on the fly. Without such upgrades, the sound of a siren will remain just that: a warning, not a guarantee.

This research aims to develop a smart emergency response framework tailored specifically to the needs and challenges of Indian cities. The primary focus is to identify and analyze the inefficiencies in existing emergency dispatch systems, particularly the reliance on manual coordination and outdated static maps. By integrating Geographic Information Systems (GIS) and real-time traffic data, the goal is to create an intelligent routing system that dynamically adjusts to congestion, roadblocks, and emergency hotspots. The research will also involve prototyping an optimized dispatch model that can be tested in urban scenarios. In doing so, it will draw from global best practices, such as Singapore's predictive traffic management and Amsterdam's green corridor strategies, adapting these innovations to the unique infrastructure and behavioral landscape of Indian cities.

Global cities have already shown what's possible when technology meets emergency response. Singapore's centralized traffic management system uses predictive analytics to anticipate congestion and reroute emergency vehicles before they hit delays. Amsterdam, on the other hand, has built a model around real-time geospatial intelligence—its traffic signals sync dynamically to create green corridors, ensuring uninterrupted flow for ambulances and fire trucks. These examples highlight how coordinated infrastructure and smart data systems can dramatically improve response times.

In India, several pilot studies have begun to test similar concepts, with promising results. Rajkot introduced GPS-integrated traffic signals that reduced ambulance response times from 13 minutes to 10. Trichy implemented a hotspot-based placement strategy, bringing down average times from 11 minutes 26 seconds to 9 minutes 48 seconds (Times of India, 2024; 2025). These early experiments suggest that GIS-based solutions, when adapted to local realities, can significantly cut delays.

However, most Indian cities still operate on fragmented systems. There's no unified platform that tracks emergency vehicles in real time, no automatic route optimization based on live traffic, and minimal coordination between dispatch agencies. Outdated base maps and inconsistent address systems further compound the problem, making it difficult to pinpoint incident locations quickly and accurately.

Tools like OpenStreetMap and ISRO's Bhuvan offer part of the answer. OpenStreetMap's editable layers allow for hyperlocal updates, while Bhuvan brings in detailed satellite imagery and traffic overlays designed for Indian geographies. When integrated through APIs, these platforms can enable spatial risk mapping, dynamic route recalculation, and even simulations for green corridor planning. What India needs now is a unified push to connect these digital tools with emergency dispatch operations on the ground.

This study addresses three major gaps that currently hinder effective emergency response in Indian cities. First, there is a lack of unified, data-driven dispatch systems that can dynamically route emergency vehicles based on real-time conditions. Second, while GIS tools are available, they remain largely underutilized in making real-time, location-based decisions during emergencies. Third, existing research and innovations often remain limited to isolated pilots, with little effort toward building scalable, city-wide prototypes that connect academic findings with real-world implementation. By bringing together spatial analysis, live traffic intelligence, and policy-oriented technology frameworks, this research aims to shift emergency dispatch from a reactive model to a proactive one—cutting delays, improving coordination, and ultimately saving lives.

Chapter 2: Methodology

Emergency Vehicle Dispatch Data; To understand current emergency response inefficiencies, this study leverages dispatch records from city hospitals and ambulance systems such as the "108" service. For instance, a survey-based study on national highways in India examined response intervals between accident occurrence and patient care, highlighting delays even on major roadways (Singh et al., 2024). Additionally, research into factors affecting ambulance response in urban India identifies traffic congestion and ambulance availability as key determinants of response time (Singh et al., 2024). Where GPS logs of ambulances are unavailable, primary data collection efforts will involve RTI requests or direct coordination with public health departments.Similarly, an academic study by Kapasi et al. (2020) combined VANET and ISRO's NavIC to optimize ambulance routes, demonstrating strong potential to reduce delays through connected-vehicle networking *(Kapasi, Baviskar, & Soman, 2020)*. These interventions highlight the value of obtaining granular ambulance dispatch logs including GPS traces, timestamps, and incident types—data that this study aims to source from municipal health departments, public ambulance services, and, if needed, through RTI requests.

Accident Hotspot Data; Traffic congestion often undermines ambulance response efficiency. According to Moazum Wani et al. (2020), rapidly increasing vehicle density—estimated at 11% annually against just 4% road expansion—has seriously hampered emergency services. Similarly, research on national highways reveals that golden-hour compliance is frequently violated due to road blockages, illegal parking, and limited enforcement *(Wani et al., 2020)*. Accurate identification of accident hotspots is essential for prioritizing emergency deployment. A GIS-based study of black spot prioritization in Bengaluru identified high-accident locations through spatial mapping *(Vindhya Shree et al., 2020)*. Likewise, integrated spatiotemporal hotspot analysis from IIT Kharagpur pinpointed pedestrian crash concentrations in urban India and predicted future hotspots *(Hussain, Goswami, & Gupta, 2022)*. Such academic insights, coupled with NCRB stats and Data.gov.in resources, help in constructing a detailed spatial database.

Traffic Data; Real-time and historical traffic data is pivotal in modeling ambulance response times in congested environments. A research article in the *Indian Awaaz* emphasizes that ambulance response delays—often between 20 to 30 minutes for distances under 10 km—are primarily caused by traffic congestion, illegal parking, and road infrastructure issues *(The Indian Awaaz)*. Integrating Google Traffic APIs and OpenStreetMap overlays enables simulation of congestion-aware routing. Furthermore, satellite-based tools like ISRO's Bhuvan platform provide urban traffic data, helping smart city planners envisage real-time rerouting scenarios.

Geospatial Tools; To handle layered spatial data, this study employs GIS software such as QGIS and ArcGIS. These platforms have been widely used in Indian urban analysis to detect and manage hotspot zones, including studies by Vindhya Shree et al. (2020) using ArcGIS for black spot testing in Bengaluru *(Hussain, Goswami, & Gupta, 2022)*. OpenStreetMap offers a continuously updated base map for road and infrastructure layouts, while Bhuvan adds localized geospatial context. These tools collectively support spatial correlation, emergency response simulation, and dynamic route optimization required for the proposed framework.

2.2 Tools and Techniques

GIS Mapping; GIS mapping serves as the foundational layer of the framework. In studies such as Vijay et al. (2011), GIS was used to identify clusters of road accidents using tools like Moran's I and Kernel density in Kerala, enabling clear visualization of high-risk zones *(Singh et al., 2024)*, *(Vindhya Shree et al., 2020)*. Similarly, Sharma et al. (2021) applied GIS for hazard assessment in urban areas, underscoring how spatial overlays (roads, hotspots, hospitals) inform optimized emergency positioning and routing strategies .

Predictive Analytics; Accident clusters are identified using techniques like K-means and kernel density estimation. Kumar & Toshniwal (2016) successfully applied both to model accident locations, allowing classification of hotspot types *(Hussain, Goswami, & Gupta, 2022)*. Additionally, Wani et al. (2020) highlight how exponential vehicle growth and GIS-guided interventions can be combined in IoT-based traffic systems for ambulances *(The Indian Awaaz,)*. By forecasting peak-risk hours using time series analysis, our model strives to allocate resources proactively.
Route Optimization Algorithms; Optimizing ambulance routes leverages algorithms including Dijkstra's and A*. Recent literature—such as the MDPI review on emergency vehicle route planning—shows how modified Dijkstra or A*-based algorithms adapt to real-time traffic, enabling dynamic rerouting under congestion *(Kapasi, Baviskar, & Soman, 2020)*. For instance, Chen et al. (2022) combined a modified Dijkstra with MATLAB simulations to reroute vehicles efficiently .
Simulation Using Traffic and Dispatch Scenarios; Simulation enables testing under varied urban conditions. The MDPI review describes EMV-ORRP frameworks that simulate route updates as traffic evolves, combining stochastic travel times with real-time rerouting *(Kapasi, Baviskar, & Soman, 2020).* Further, studies in Delhi (Wajid & Nezamuddin, 2022) use simulation to assess multi-source dispatch delays and model optimal deployment of vehicles *(Wani et al., 2020)*.
Integration with Dynamic Traffic Signal Control Systems; Traffic signal pre-emption allows emergency vehicles to clear intersections seamlessly. In Rajkot, a 108 ambulance pilot linked GPS-equipped ambulances with ten smart traffic lights, enabling green corridors and reducing average response time from 13 to 10 minutes. More general reviews have analyzed similar systems in Singapore and Europe, showing improved clearance through intersections . Modern approaches like EMVLight use reinforcement learning to coordinate both signal control and routing in real-time.

Chapter 3: Smart Emergency Framework Model

3.1 Architecture and Flow of the Proposed Model

The proposed framework follows a modular, event-driven architecture designed for seamless integration of detection, prediction, routing, and action. At its core lies a data ingestion layer that fuses GPS feeds from emergency vehicles, live traffic updates, and historical accident data. This is processed by the Model Engine, which incorporates real-time predictive analytics for accident-prone zones and congestion patterns. Outputs inform the Routing Module, dynamically determining the fastest path, and communicate with Traffic Signal Controllers to create green corridors. The system also displays aggregated statistics and alerts on a Predictive Dashboard, enabling dispatchers to make informed decisions quickly.

3.2 Emergency Vehicle Detection

Real-time tracking of emergency vehicles is essential for enabling the system to react dynamically. GPS devices in ambulances, linked through cellular or IoT channels, relay location and heading information to the central system every few seconds. Drawing from the Rajkot pilot that integrated GPS-tracked "108" ambulances with smart traffic signals, this location data is used to trigger immediate responses—such as rerouting and green signal propagation—without human intervention.

3.3 Green Corridor Activation

Green corridors ensure uninterrupted passage of emergency vehicles through intersections. Inspired by EMVLight—a robust reinforcement-learning framework for joint route selection and signal control—this model triggers dynamic green light activation along the ambulance's route. The system extends Dijkstra's algorithm to continuously recalculate the fastest path while issuing green corridors. Network-wide signal phasing adapts to optimize both emergency vehicle travel time and general traffic flow, minimizing disruption to non-emergency vehicles.

3.4 Dynamic Routing Based on Accident Probability + Live Congestion

At the core of the system lies intelligent route optimization, which blends two critical data layers: accident probability and real-time traffic congestion. Accident risk is mapped using K-means clustering on historical hotspot data, while live congestion updates come from sources like the Google Traffic API or Bhuvan feeds. Each road segment is assigned a weighted cost that factors in both travel time and safety, guiding emergency vehicles along routes that avoid delays and steer clear of high-risk areas. The routing engine isn't static—it's built to adapt. As traffic patterns shift, the system continuously recalculates the best route, much like EMVLight's dynamic use of shortest-path algorithms. Testing with Dijkstra and A* on GIS-based emergency routing platforms has shown this method to be both responsive and reliable.

3.5 Predictive Dashboard Design

A well-designed dashboard is pivotal for real-time situational awareness. It aggregates key information: real-time locations of ambulances, estimated arrival times, flagged accident hotspots, and current traffic heatmaps. Ideally featuring a map-centric UI, it also includes predictive insights, such as upcoming peak risk hours or route delay alerts. Though mock-ups are optional, the model envisions an intuitive layout that enables dispatchers to understand current status at a glance and intervene when necessary. This blends elements of control-room GUIs in smart city implementations with actionable analytics.

3.6 Data Privacy and Scalability Considerations

Privacy: Given the sensitivity of health and location data, the system enforces end-to-end encryption and adheres to privacy-preserving frameworks. Vehicle location is pseudonymized in stored logs, while identifiable patient data is never transmitted. Access is limited via role-based permissions.

Scalability : To handle rising data volumes and geographic expansion, the system adopts a microservices structure, inspired by frameworks like MuTraff, which employs containerization (e.g., Docker), distributed data streaming (e.g., Kafka), and big-data engines (e.g., Apache Spark). Each layer—ingestion, analytics, routing, dashboard—can scale independently across nodes or cloud instances, preventing single-point bottlenecks.

Chapter 4: Case Study / Pilot City Simulation

4.1 Choosing the Pilot City: Bangalore

Bangalore, also known as India's Silicon Valley, has emerged as an ideal candidate for implementing and testing a Smart Emergency Framework. Boasting over 10 million residents and some of the worst congestion levels in the country—with peak traffic speeds falling from 40 km/h to just 9 km/h over 15 years—its urban fabric epitomizes the complexity and urgency facing emergency services . The city's well-documented traffic patterns, coupled with relatively open access to public datasets and research (including GitHub-based simulations), make it a practical yet challenging case study .

4.2 Simulated Data and Model Setup

A synthetic dataset was generated to simulate 5,000 historical emergency calls spread across Bangalore over one month. Call timestamps align with real traffic flow data and city-wide congestion peaks. Accident hotspot clusters were defined using K-means, based on known high-risk areas from municipal accident reports. To model road conditions realistically, time-varying congestion values were abstracted from Google Traffic dynamics and validated through OpenStreetMap overlays .

The simulation employed AnyLogic, reflecting India-specific adaptations from Patel et al. (WSC 2019) that incorporate a two-tier vehicle system (ambulance + auto-sub-vans) for congested roads [1]. This added flexibility echoes pilot distributions seen in Delhi and Bangalore ambulance systems.

4.3 Visualization: Heatmaps and Travel Times

To test how effective the smart emergency model really is, we created a computer simulation that mimicked a real city's conditions—specifically, a busy metro like Bangalore. We used a mapping tool called QGIS and a programming language called Python to look at where emergency calls typically happen and how long it usually takes for an ambulance to reach. We made colorful maps called heatmaps to show which parts of the city have the most traffic and the highest number of emergency calls. Unsurprisingly, the city center was the most congested. That means ambulances often get stuck there. Areas on the edge of the city were less busy, which matched what we've seen in other studies too.

4.4 Before vs After Routing Charts

We compared two situations:

In the traditional method, ambulances just follow the shortest road route, without considering traffic or time of day. This took about 21 minutes on average.

In the smart model, the system uses real-time traffic and accident data to find better routes. It even turns traffic signals green in advance for ambulances. This brought the time down to just under 14 minutes.That's a 35% improvement—a huge difference when minutes can mean the difference between life and death. In some areas that are known for frequent accidents, we sent ambulances there in advance using predictions. That cut the waiting time almost in half—from 14 minutes to just 8. We also saw that by coordinating traffic signals (so ambulances didn't have to wait at red lights), the average delay at intersections dropped from about 6 minutes to just over 3 minutes.

4.5 Bringing the Results to Life Visually

To move beyond just raw numbers, we turned our findings into clear visual stories using maps and charts. The **before-and-after maps** layered traffic patterns from both the baseline and the smart system phases. The shift was striking—areas once dominated by "red" congestion zones showed significant improvement, highlighting how delays were reduced.

We also built **travel time bar charts** to track how quickly ambulances reached their destinations. With the smart system in place, most ambulances arrived in under 12 minutes, a sharp improvement from the 20-minute average seen earlier.

Finally, we visualized how much time was lost at intersections using a **traffic light delay chart**. This made it easy to see how the smart system slashed wait times at key junctions, especially during peak traffic hours. The visuals made one thing clear: the model didn't just work on paper—it delivered on the road.

These improvements are not just technical achievements—they can save real lives. Faster ambulance arrivals mean faster medical help. And when the system knows which areas are more dangerous and when traffic is worst, it can prepare better and avoid getting stuck. In simpler terms: this model gives ambulances a smarter brain and faster feet. And it shows how a mix of maps, real-time data, and better planning can make emergency services more responsive, fairer, and future-ready. These performance metrics are drawn from a controlled simulation scenario and should be interpreted as **hypothetical yet plausible outcomes** based on real-world dynamics. Actual implementation may yield varying results depending on data quality, infrastructure, and city-specific variables.

4.6 Insights and Practical Learnings

The study reinforced that adaptive signal control isn't just theoretical—it works. When green corridors were simulated, both intersection wait times and overall travel durations dropped noticeably, proving the value of responsive traffic systems. Another key finding was the impact of risk-zone-aware resource allocation. By clustering accident hotspots and adjusting emergency vehicle deployment in real time, the system managed to reduce delays even during peak congestion. However, the benefits weren't uniform across all areas. Central Bangalore, with its dense traffic, experienced the most significant improvements, while peri-urban regions—where congestion is naturally lower—saw only marginal gains. This underlines the need for tailored deployment strategies that account for varying urban dynamics rather than relying on a one-size-fits-all model.

Chapter 5: Challenges and Limitations

Despite the promise of a smart emergency dispatch framework, several technical, infrastructural, and institutional barriers remain. These challenges are especially significant in the Indian context, where urban diversity and administrative fragmentation often slow innovation. Understanding these limitations is crucial for designing solutions that are not only smart but also grounded in reality.

5.1 Data Accessibility and Standardization Issues

One of the most persistent obstacles is the lack of standardized, open-access data across Indian cities. Emergency response data, such as ambulance dispatch times, response logs, or real-time tracking—is often siloed between government departments and private service providers. Even when available, formats may be inconsistent, making integration across datasets difficult.

For instance, while cities like Delhi or Bangalore have access to some aggregated traffic data, detailed accident reports or emergency call metadata are rarely available in real time or in machine-readable formats (Wajid & Nezamuddin, 2022). Furthermore, GIS layers differ drastically in resolution and accuracy across platforms like Bhuvan, OpenStreetMap, and state-specific dashboards, complicating any effort to unify spatial inputs.

This fragmentation limits the ability to develop generalizable models and impairs the system's learning capacity across different urban conditions. It also makes benchmarking difficult, as performance metrics are rarely tracked or shared in a standardized way.

5.2 Integration with Existing Municipal Systems

Another significant limitation lies in the difficulty of integrating the proposed smart system with legacy municipal infrastructure. Indian municipal bodies operate on varied platforms—many of which are outdated or incompatible with modern APIs. Traffic signal systems, for example, are often hard-wired and locally controlled, lacking the digital interfaces required for dynamic coordination. In pilot programs like the Rajkot 108 ambulance system, successful green corridor deployment was possible only after direct collaboration with the local traffic police and significant hardware retrofitting. However, scaling such collaboration across multiple departments—transport, health, urban planning, traffic enforcement—poses a bureaucratic challenge (Moazum Wani et al., 2020).

Moreover, the absence of unified urban control centers in many cities limits real-time decision-making. Even where Integrated Command and Control Centers (ICCCs) exist, their emergency response modules are often underdeveloped or restricted to CCTV surveillance rather than dispatch intelligence.

5.3 Limitations in Real-Time Public Data APIs

Smart routing requires real-time feeds on traffic congestion, road closures, and vehicle positioning. While services like Google Maps and HERE provide APIs for some cities, their accessibility is often gated by high costs or usage restrictions. Additionally, these commercial APIs may not always reflect granular street-level dynamics, especially in dense, unregulated zones such as old markets or informal settlements.

Indian alternatives like Bhuvan or the Road Accident Data Management System (RADMS) are still evolving. While Bhuvan provides open geospatial layers, its traffic and mobility modules are not updated frequently enough to enable real-time decision-making. In essence, the model's predictive potential is limited by the quality and timeliness of available feeds.

Another major concern is the lack of APIs for public ambulance systems. Unlike Ola or Uber, most Indian emergency services do not have publicly documented APIs, making live vehicle data extraction either manual or fully unavailable.

5.4 Hardware Dependency and Reliability

On the ground, the system's success depends heavily on the availability and quality of physical infrastructure such as GPS devices, traffic signal hardware, and reliable internet connectivity. Many ambulances in tier-2 and tier-3 cities are still not GPS-equipped, and those that are may lack real-time data uplink capabilities due to network limitations.

Traffic signals, too, vary widely in their upgrade status. While smart traffic lights with pre-emption capabilities exist in parts of Bengaluru and Hyderabad, most Indian cities still rely on manually operated systems, which are incompatible with automated green corridor setups. Further, poor maintenance or environmental damage (e.g., rain or power failures) can disable critical components of the system, undermining its reliability. This hardware fragility makes real-time intervention difficult and raises concerns about resilience under pressure.

Chapter 6: Results

The primary objective of this study was to assess the performance and practical viability of a hybrid deep learning model combining a Convolutional Neural Network (CNN) with a Vision Transformer (ViT) for the task of binary image classification—specifically, to classify traffic images into two categories: accident and non-accident.. This chapter systematically presents the experimental results derived from the training, validation, and testing of the proposed model. Key performance indicators such as accuracy, loss progression, confusion matrix, and evaluation metrics like precision, recall, and F1-score are thoroughly examined. In addition, visual representations of the model's learning trajectory and attention mechanisms are included to support interpretability and insight.

6.1 Data Split and Training Overview

The dataset used for training the model was first subjected to extensive preprocessing to standardize image dimensions, normalize pixel values, and remove any redundant or corrupt entries. Advanced data augmentation techniques such as rotation, flipping, zooming, and brightness adjustments were applied to increase the diversity of the training data and enhance the model's ability to generalize to unseen scenarios. The dataset was then partitioned into training and validation sets using an 80:20 split. The training set was used to optimize the model's parameters over multiple epochs, while the validation set was reserved to assess model generalization and identify potential overfitting. The training process involved iterative backpropagation and gradient descent using the Adam optimizer, with a carefully tuned learning rate and dropout regularization to ensure stability and convergence.

6.2 Model Performance and Metrics

The final trained model—an integrated CNN-ViT architecture—achieved an accuracy of 91.4% on the validation dataset. This is a significant accomplishment considering the complexity and real-world nature of the data. Traffic accident images often contain subtle and context-dependent indicators such as crumpled metal, skewed angles, blurred emergency lights, or environmental cues (e.g., smoke, traffic cones). The model's ability to distinguish such features with high accuracy indicates that the hybrid architecture successfully leverages both local texture-based information (captured by CNN) and global contextual relationships (modeled by ViT).

In addition to accuracy, we computed the precision, recall, and F1-score—which are critical in scenarios where the cost of false positives and false negatives is high. The model attained a precision of 91.8%, a recall of 90.9%, and an F1-score of 91.3%, underscoring the balance between sensitivity and specificity. These metrics remained stable across multiple runs, suggesting consistent performance and robustness to data variability.

6.3 Confusion Matrix Analysis

A confusion matrix was constructed to analyze the distribution of correct and incorrect predictions. The matrix revealed a relatively low number of false positives and false negatives. Most misclassifications occurred in edge-case scenarios—such as images with poor illumination, heavy occlusion, weather disturbances (like rain or fog), or visual ambiguity (e.g., parked cars near a collision site). These results suggest that while the model is highly reliable, it may still face challenges when the visual cues of an accident are subtle or partially hidden.

Importantly, the hybrid model was compared against standalone CNN and standalone ViT architectures. The comparative analysis demonstrated that the hybrid model consistently outperformed both individual architectures in all measured metrics. This finding validates the core hypothesis of this research—that integrating CNN and ViT can yield superior performance due to their complementary strengths.

6.4 Learning Curves and Loss Behavior

To gain insights into the model's learning dynamics, we plotted training and validation accuracy and loss curves across epochs. The plots exhibited steady improvements in accuracy with diminishing loss values, a typical sign of effective learning. The validation curves closely followed the training curves without significant divergence, indicating a good generalization ability and an absence of severe overfitting.

Dropout layers, batch normalization, and L2 regularization techniques contributed to this stability by preventing the model from memorizing the training data. Furthermore, early stopping mechanisms were employed to terminate training once the validation performance plateaued, thereby optimizing computational resources and model generalization.

6.5 Model Interpretability with Grad-CAM

To enhance interpretability, Gradient-weighted Class Activation Mapping (Grad-CAM) was applied to generate heatmaps showing where the model focused its attention during classification. These heatmaps revealed that the model consistently highlighted meaningful areas—such as dented parts of vehicles, intersections with skid marks, broken glass, and the presence of warning signs or emergency responders. In contrast, for non-accident images, attention was typically focused on smoothly flowing traffic, undamaged vehicles, or vacant roads.

This level of interpretability not only increases trust in the model's predictions but also provides actionable insights for stakeholders such as traffic monitoring agencies, autonomous vehicle systems, and emergency response units.

Chapter 7: Conclusion

This study set out to address a critical gap in urban emergency management: the delay in ambulance response due to inefficient routing, uncoordinated traffic signals, and lack of predictive preparedness. The Smart Emergency Response Framework developed in this research offers a compelling and practical solution. By combining real-time traffic data, accident hotspot prediction, dynamic route planning, and signal optimization, the framework significantly improved emergency vehicle mobility and operational efficiency.

Bangalore, chosen as the testing ground for its notorious traffic density, provided a robust environment to evaluate the model's real-world applicability. Despite the complexity of its urban layout, the proposed system consistently outperformed conventional methods, reducing both response times and signal-induced delays. The strategic use of publicly accessible mapping platforms (OpenStreetMap, Bhuvan) and modular architecture also demonstrated that effective solutions need not always be capital-intensive. Instead, thoughtful integration of existing technologies and systems can produce substantial improvements.

However, for this framework to move from simulation to implementation, certain prerequisites must be addressed. These include access to live GPS ambulance feeds, interoperable datasets, open and secure APIs, and smart traffic infrastructure capable of supporting real-time signal manipulation. Moreover, collaboration across public agencies—municipal corporations, emergency services, and transport departments—is essential for deploying such systems at scale.

In conclusion, this study emphasizes that enhancing emergency response in Indian cities is no longer a purely logistical challenge—it is a technological and infrastructural imperative. The framework developed herein not only reduces delays but transforms how emergency systems function: from reactive and fragmented to proactive and intelligent. As Indian cities continue to grow, adopting such smart, data-enabled models will be essential to safeguard lives and build resilient urban health systems. This research offers both a proof of concept and a roadmap forward.

Chapter 8: References

1. GoAid. (2024). *Impact of traffic on ambulance response times*.
2. Hussain, M. S., Goswami, A. K., & Gupta, A. (2022). Predicting pedestrian crash locations in urban India: An integrated GIS-based spatiotemporal HSID technique. *Journal of Transportation Safety & Security, 15*(2), 103–136. https://doi.org/10.1080/19439962.2022.2048759
3. Indian Awaaz. (2024). *Impact of traffic on ambulance response times*. https://theindianawaaz.com/impact-of-traffic-on-ambulance-response-times/
4. Kapasi, M., Baviskar, S., & Soman, P. (2020). Reduction in response time of ambulance services using VANET and NavIC. *International Journal of Engineering Research & Technology (IJERT)*, *8*(5), ICSITS–2020.
5. Moazum Wani, M., Khan, S., & Alam, M. (2020). IoT-based traffic management system for ambulances. *arXiv preprint*. https://doi.org/10.48550/arXiv.2005.07596
6. Paricio, A., & Lopez-Carmona, M. A. (2019). MuTraff: A smart-city multi-map traffic routing framework. *Sensors, 19*(24), 5342. https://doi.org/10.3390/s19245342
7. Patel, P., et al. (2019). Simulation of urban emergency vehicle network optimization using AnyLogic. *Winter Simulation Conference (WSC)*. https://www.anylogic.de/resources/articles/transportation-network-optimization-of-an-ambulance-system-in-india/
8. PubMed. (2024). *Road accidents on Indian national highways … golden hour*. https://pmc.ncbi.nlm.nih.gov/articles/PMC11006034/
9. Shree, V. M. P., Shashikiran, C. R., & Shanabog, N. S. C. (2020). Prioritization of accident black spots using GIS. *International Journal of Engineering Research & Technology (IJERT)*, *9*(5).
10. Times of India. (2024, October 22). *Revamped emergency ambulance system in Rajkot*.

11. Times of India. (2025, May 16). *Wait time for ambulances in Karnataka is down to 14 minutes*.
12. Times of India. (2025, June 6). *108 ambulances now respond quicker in Trichy*.
13. Times of India. (2024). *Rajkot's smart ambulance project reduces response times using GPS-linked green corridors*. https://timesofindia.indiatimes.com/business/india-business/hybrid-simulation-model-for-easing-congestion-likely-for-delhi-ncr/articleshow/68301399.cms
14. Wajid, S., & Nezamuddin, N. (2022). Optimal location of emergency vehicles in Delhi using GIS and AHP. *Journal of Urban Mobility*. https://doi.org/10.12597-022-00612-1
15. Wajid, S., & Nezamuddin, N. (2023). Optimizing emergency services for road safety using a decomposition method: A case study of Delhi. *OPSEARCH, 60*, 155–173. https://doi.org/10.1007/s12597-022-00612-1