Multidisciplinary Journal of Research in Engineering and Technology
ISSN: 2348-6953
Volume 13 Issue 2s (2026) | Pages: 96-108

A Hybrid Multi-Agent Emergency Response Simulation Platform with Centralized Dispatch and Localized Agent Autonomy

Mayuri Fegade[1], Darshan Ingale[2], Narayani Shelke[3], Prateek Bodre[4], Prathmesh Nandgaonkar[5]

[1,2,3,4,5]Department of AI & DS, Dr. D. Y. Patil College of Engineering & Innovation, Pune, India

Peer Review Information

Type: *Article*
Received: *26 March 2026*
Revised: *10 April 2026*
Accepted: *24 May 2026*
Published: *15 June 2026*

Abstract

Urban emergency response systems are complex, high-risk, and hard to evaluate directly in real-world settings due to operational limits, costs, and safety concerns. Just as circuit designers use virtual prototyping platforms like Tinkercad before deploying physical hardware, this work introduces a simulation-first approach for planning and evaluating emergency responses. We develop a hybrid multi-agent simulation platform that allows for "test-before-deploy" experimentation for ambulance dispatch and urban incident management policies. The proposed system is built on a realistic urban road network derived from OpenStreetMap and runs using SUMO (Simulation of Urban Mobility) with TraCI integration. It models civilian traffic, hospital infrastructure, random incident generation, and a fleet of ambulances acting as autonomous agents. Each ambulance keeps localized state, intent, and behavior policies, allowing for decentralized decision-making for patrol and response. A centralized dispatch layer coordinates incident assignments using a scoring system, which can be enhanced with a tabular reinforcement learning policy for selecting primary responders.

The platform adds more operational realism through severity-based responder allocation, scene dwell times, hospital handover delays, hotspot-aware patrol rebalancing, and localized traffic-signal preemption to simulate green corridors. An expandable experimental layer allows for synthetic modeling of weather conditions, traffic congestion patterns over time, variations in district-level demand, and hospital operational limits. The system provides structured operational metrics, including response times, completion rates, and resource use, which are analyzed through a dashboard-driven process. This setup allows for systematic comparison of dispatch strategies, fleet configurations, and environmental conditions under controlled, reproducible scenarios.

The main contribution of this work is a modular and rollback-safe simulation platform that connects heuristic multi-agent control with learning-based dispatch in a realistic urban environment. By allowing rapid prototyping and evaluation of emergency response strategies before real-world deployment, the platform acts as a digital testbed for research in smart-city mobility, multi-agent coordination, and optimizing emergency logistics.

Keywords: Emergency Response Systems; Multi-Agent Simulation; SUMO; Ambulance Dispatch; Reinforcement Learning; Smart City Mobility.

How to Cite This Article

Fegade, M., Ingale, D., Shelke, N., Bodre, P., & Nandgaonkar, P. (2026). A hybrid multi-agent emergency response simulation platform with centralized dispatch and localized agent autonomy. *Multidisciplinary Journal of Research in Engineering and Technology*, 13(2s), 96–108.

Introduction

A rapid and effective emergency medical response is essential for urban infrastructure. In densely populated cities like Pune, delays in ambulance dispatch, routing, and hospital handover can greatly affect patient outcomes. However, designing and assessing better response strategies in real-world settings is tough due to safety risks, operational costs, limited control, and the challenge of conducting repeatable experiments under similar conditions.

To overcome these challenges, simulation-driven methods have come up as a practical way to study urban mobility and emergency logistics. Tools like SUMO (Simulation of Urban MObility), along with programmatic control via TraCI, allow for modeling realistic road networks, traffic flows, and vehicle behavior. These platforms let us evaluate system-level interventions, such as dispatch strategies, fleet sizing, and traffic-signal prioritization, without disrupting actual operations.

This work follows a “simulation-before-deployment” approach, similar to how engineers use virtual prototyping environments like Tinkercad in electronics design. Just as circuits are validated in simulation before hardware implementation, we propose a digital testbed that enables the design, testing, and refinement of emergency response policies before applying them in real life.

This approach reduces risk, enhances reproducibility, and allows for systematic exploration of complex operational trade-offs. We present a hybrid multi-agent emergency response simulation platform that combines centralized dispatch coordination with localized agent autonomy. The system models a fleet of ambulances operating over an urban road network derived from OpenStreetMap, along with civilian traffic and hospital infrastructure.

Each ambulance acts as a stateful agent with its own intent, memory, and behavioral policies, allowing for decentralized decision-making in patrol and incident response. Meanwhile, a centralized dispatcher assigns ambulances to incidents using a scoring system, which can be improved with a reinforcement learning component for selecting primary responders.

The platform includes several operational features to enhance realism, such as severity-based incident modeling, multi-ambulance coordination, hospital-aware routing, scene dwell and handover delays, and localized traffic-signal preemption to simulate green corridors.

There is also an extensible experimental layer that supports synthetic modeling of factors like weather conditions, changing congestion levels, district-level demand variation, and hospital capacity limits.

The main goal of this work is not to claim optimal dispatch performance but to provide a flexible and reproducible environment for testing emergency response strategies under controlled conditions. By combining rule-based multi-agent coordination with an optional learning-based dispatch layer, the platform allows for comparative analysis across various operational scenarios.

The contributions of this paper are threefold:

- The design of a SUMO-based emergency response simulation framework based on a real urban road network.
- A hybrid control architecture that blends centralized dispatch with decentralized agent behavior.
- A modular experimental layer and KPI-driven evaluation process for systematic analysis of emergency response performance.

The rest of this paper is organized as follows: Section 2 reviews related work in emergency response simulation and multi-agent systems; Section 3 describes the system architecture and modeling approach; Section 4 details the dispatch and control mechanisms; presents the experimental setup and evaluation metrics; Section 5 discusses results and insights; and Section 6 concludes with limitations and future work.

Related Work

Research on emergency response systems covers various areas, including traffic simulation, dispatch optimization, and multi-agent coordination. This section reviews the most relevant literature and positions the proposed system within this context.

Traffic Simulation for Urban Mobility

Microscopic traffic simulators like SUMO (Simulation of Urban Mobility) have been widely used to study urban transportation systems under controlled and reproducible conditions. These tools allow for detailed modeling of vehicle interactions, road networks derived from OpenStreetMap, and traffic control methods. The availability of interfaces like TraCI enables external controllers to influence simulation behavior dynamically, making SUMO especially useful for integrating custom logic such as emergency vehicle routing and signal pre-emption.

Previous work has used SUMO to analyze congestion patterns, evaluate intelligent transportation systems, and simulate emergency vehicle prioritization. However, many studies focus either on traffic-level interventions, such as optimizing signal timing, or on isolated routing strategies. They often neglect modeling the entire lifecycle of emergency response, which includes dispatch, on-scene handling, and hospital transport.

Emergency Dispatch and Optimization Approaches
Traditional emergency dispatch systems are often framed as optimization problems, including variations of the vehicle routing problem (VRP) and dynamic assignment models. Strategies such as nearest-vehicle dispatch, greedy heuristics, and integer programming aim to minimize response time and maximize coverage. More complex methods involve modeling uncertain demand and dynamic repositioning to deal with unpredictability in incident occurrence. While these methods provide strong theoretical guarantees under simpler assumptions, they often face difficulties in scaling or adjusting to real-world complexities like traffic variability, differences in incident severity, and hospital congestion.

In contrast, heuristic-based dispatch systems are frequently used in practice due to their efficiency and adaptability. The system proposed here follows a similar philosophy, using a scoring-based heuristic to rank ambulance-incident assignments while taking into account several operational factors such as severity, travel time, and coverage balance.

Multi-Agent Systems in Emergency Response
The use of Multi-Agent Systems has gained interest for modeling decentralized coordination in complex settings. In these systems, individual agents operate based on their local state and policies while working together toward shared objectives. Previous studies have treated emergency vehicles as agents capable of autonomous routing, negotiation, or cooperative task allocation. These approaches are beneficial for capturing emergent behavior and scalability in large urban areas. However, fully decentralized systems can have trouble maintaining global coordination and consistency, especially in high-stakes fields like emergency response.

The approach in this work adopts a hybrid structure, combining centralized dispatch decision-making with localized agent autonomy. This setup allows a balance between global coordination through dispatch and adaptive local behavior via agent-level patrol and response logic.

Reinforcement Learning for Dispatch and Control
Recent advances in Reinforcement Learning have led to its use in transportation and logistics problems, such as fleet management and dynamic dispatch. Deep reinforcement learning techniques have been applied to learn policies for taxi dispatch, ride-sharing, and emergency response amid uncertain demand. However, many RL-based methods require extensive training data, complex state representations, and significant computational power. Moreover, purely learning-based systems may lack interpretability and stability in safety-critical situations.

In this work, reinforcement learning is used in a limited and controlled way through a tabular Q-learning component for primary ambulance dispatch. Instead of replacing the entire control logic, the learning module operates alongside a heuristic framework, allowing for gradual experimentation while keeping the system stable.

Traffic Signal Preemption and Green Corridors
Traffic signal preemption for emergency vehicles, often called green corridors, has been studied as a way to reduce travel time and improve response efficiency. Existing methods range from rule-based signal overrides to coordinated corridor-level optimization.

Most previous implementations focus on either static priority rules or centralized optimization of signal timing. In contrast, the current system employs a localized, short-horizon preemption strategy, where traffic lights ahead of an ambulance are temporarily switched to green. While simplified, this method fits naturally within the simulation environment and allows for the evaluation of signal priority effects without the need for full network-wide coordination.

Summary and Positioning
Existing work generally addresses separate aspects of emergency response, such as routing, dispatch optimization, or traffic control. Few systems combine all major components—incident generation, multi-ambulance coordination, hospital-aware routing, and signal preemption—within a single, flexible simulation framework. The proposed platform contributes to this area by combining:

- a realistic urban traffic simulation environment,
- a hybrid centralized–decentralized control structure,
- an optional reinforcement learning dispatch component,
- and a structured metrics pipeline for evaluation.

Methodology

This section describes the design and implementation of the proposed hybrid multi-agent emergency response simulation platform. The system integrates traffic simulation, centralized dispatch, and decentralized agent behavior inside a unified control loop, allowing controlled experimentation of emergency response strategies.

Simulation Environment

The simulation environment is built using SUMO (Simulation of Urban Mobility), with runtime interaction enabled through TraCI. The road network is based on OpenStreetMap data, representing an urban area of Pune. This ensures that the simulation operates on a geographically realistic and structurally accurate road topology.

Civilian traffic is generated using stochastic trip generation, creating a continuous flow of passenger vehicles. Hospitals are sourced from map metadata (e.g., amenity and healthcare tags) and mapped onto valid drivable edges to ensure accessibility within the network.

The simulation runs in discrete time steps, where each step corresponds to roughly one second of real-world time. At every step, the external controller interacts with the simulator to:

- advance traffic dynamics,
- update vehicle states,
- generate new incidents,
- and execute dispatch and control decisions.

This close interaction between simulation and control logic enables real-time intervention and adaptive system behavior.

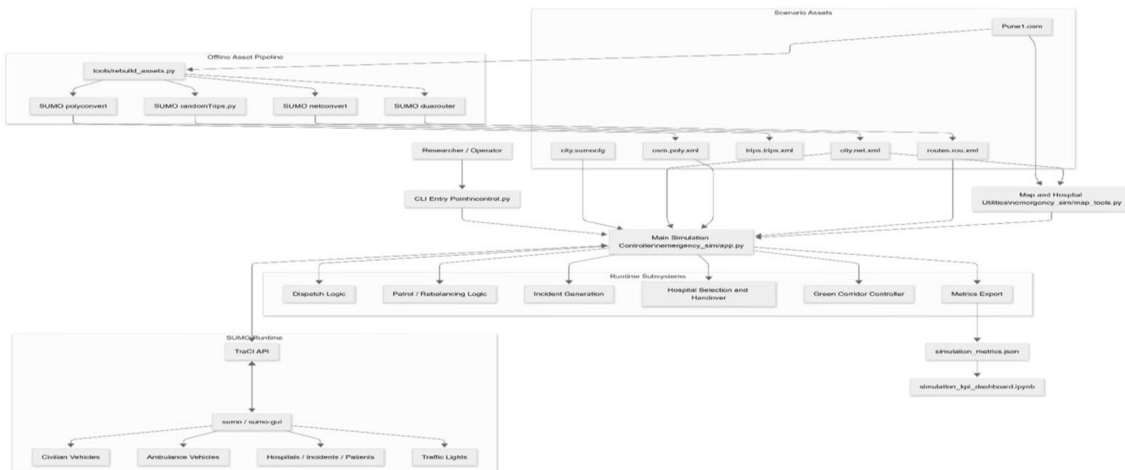

Fig. 1. *System Architecture Overview*

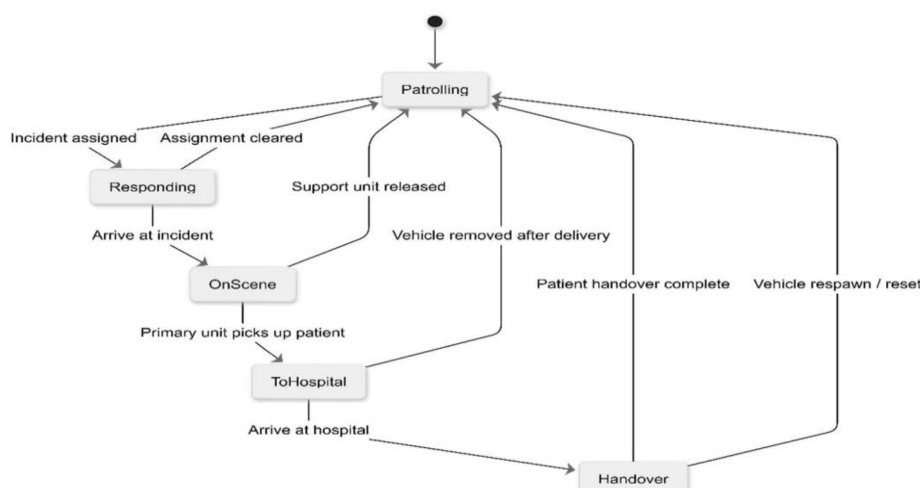

Fig. 2. *Ambulance Mission State Machine*

System Architecture

The system follows a hybrid layered architecture made up of three primary components:

- Centralized Dispatch Layer: This layer is responsible for assigning ambulances to incidents using a global view of the system. It evaluates available resources and prioritizes assignments based on operational criteria.
- Decentralized Agent Layer: Each ambulance acts as an independent agent with its own state, intent, and behavior logic. Agents make local decisions regarding navigation, patrol, and response execution.

- Environment Layer: This layer represents the simulated urban system, including traffic flow, road networks, incidents, and hospital infrastructure.

This architecture balances global coordination and local autonomy. The centralized dispatcher ensures efficient resource allocation, while decentralized agents enable flexible and context-aware behavior during execution. This hybrid design avoids the limitations of purely centralized or fully decentralized systems.

Incident Generation Model

Incidents are generated dynamically to simulate emergency demand within the urban environment. A hybrid generation strategy is used to improve realism:

1. Vehicle-based incidents: Active civilian vehicles are probabilistically selected and converted into incident sources, anchoring incidents to realistic traffic conditions.
2. Random spatial incidents: If no suitable vehicle is available, incidents are generated at randomly selected drivable edges.

Each incident is assigned a severity level $(s \in \{1,2,3\})$, denoting increasing levels of urgency. Severity is sampled randomly to reflect varying emergency conditions.

The number of required responders is calculated as:

$$R = \min(N, \max(1, b - 1) + s)$$

where **N** is the total ambulance fleet size and **b** is a base responder parameter. This formulation ensures that higher-severity incidents receive more resources while respecting fleet constraints.

This model captures both spatial and temporal variability in emergency demand while maintaining controllability for experimentation.

Dispatch Model

Dispatch decisions rely on a scoring-based heuristic that assesses the suitability of each ambulance for every active incident. For each ambulance-incident pair, a composite score is determined based on various operational factors:

- Incident severity,
- time since incident creation (age),
- estimated travel time (eta),
- hotspot intensity (historical incident density),
- ambulance experience (mission history),
- coverage balance (spatial distribution of resources).

The scoring function is defined as:

$$score = role\_bias + 52s + 0.65a + 10h + e - 0.85\eta - c$$

where:

- **s** : severity
- **a** : incident age
- **h** : hotspot score
- **e** : experience bonus
- **η** : estimated travel time
- **c** : coverage penalty

Assignments are performed greedily in descending order of score. The system enforces constraints such as:

- A single primary responder per incident,
- Additional support responders based on severity,
- And one active assignment per ambulance.

This approach allows for quick, real-time dispatch decisions while considering multiple contextual factors. Although heuristic in nature, the model offers flexibility and extensibility for incorporating learning-based components in the future.

Multi-Agent Ambulance Model

Each ambulance is designed as an autonomous agent working within a decentralized decision framework. Every agent maintains:

- A current operational state,
- an assigned incident (if applicable),
- a target hospital,
- a short-term memory of past actions and decisions,
- and an intent representation guiding its behavior.

The behavior of each ambulance is directed by a finite-state machine:

patrolling → responding → on_scene → to_hospital → handover → patrolling

In the patrolling state, ambulances reposition themselves based on hotspot memory, coverage gaps, and exploration strategies. When assigned to an incident, the agent moves to a response state, heading toward the incident location. Upon arrival, the agent enters the on-scene state for stabilization, then transitions to the hospital for patient transport, and finally completes the handover before returning to patrol.

This design allows for decentralized execution while maintaining consistency with centralized dispatch decisions, enabling agents to respond locally to changing traffic and environmental conditions.

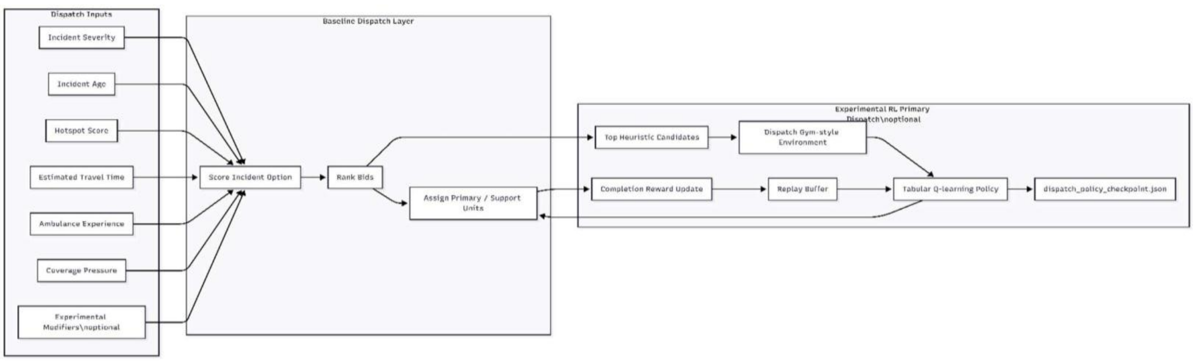

Fig. 3. *Dispatch and Decision Architecture*

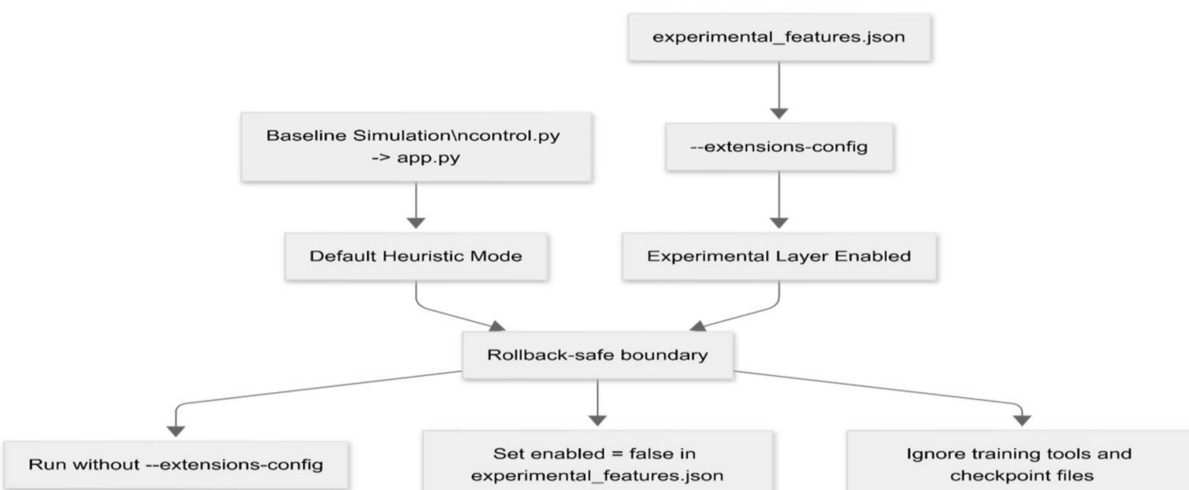

Fig. 4. *Rollback Boundary*

Hospital and Transport Model

Hospitals are selected dynamically based on the feasibility of routes and estimated travel times calculated through the simulation environment. Each incident has a designated hospital at the time of patient pickup.

The transport process consists of three phases:

- On-scene stabilization: a fixed dwell time for initial emergency care,
- Transport phase: movement of the ambulance from incident location to hospital,
- Handover phase: a delay representing patient transfer to hospital staff.

This modeling captures the full lifecycle of emergency response rather than focusing only on dispatch and routing. In extended configurations, hospital selection can factor in additional operational constraints like traffic congestion, queue lengths, and resource availability.

Traffic Signal Preemption
To simulate emergency vehicle prioritization, a localized traffic signal preemption mechanism is put in place. When an ambulance is responding or transporting a patient:

- Traffic signals within a defined distance ahead of the vehicle are identified,
- corresponding signal phases are temporarily overridden to allow a green light for passage,
- original signal programs are restored after a short brief duration.

This mechanism simulates “green corridors” to improve travel efficiency for emergency vehicles without needing complete network-wide signal coordination. The approach remains lightweight while capturing key benefits of signal priority.

Experimental Learning Component

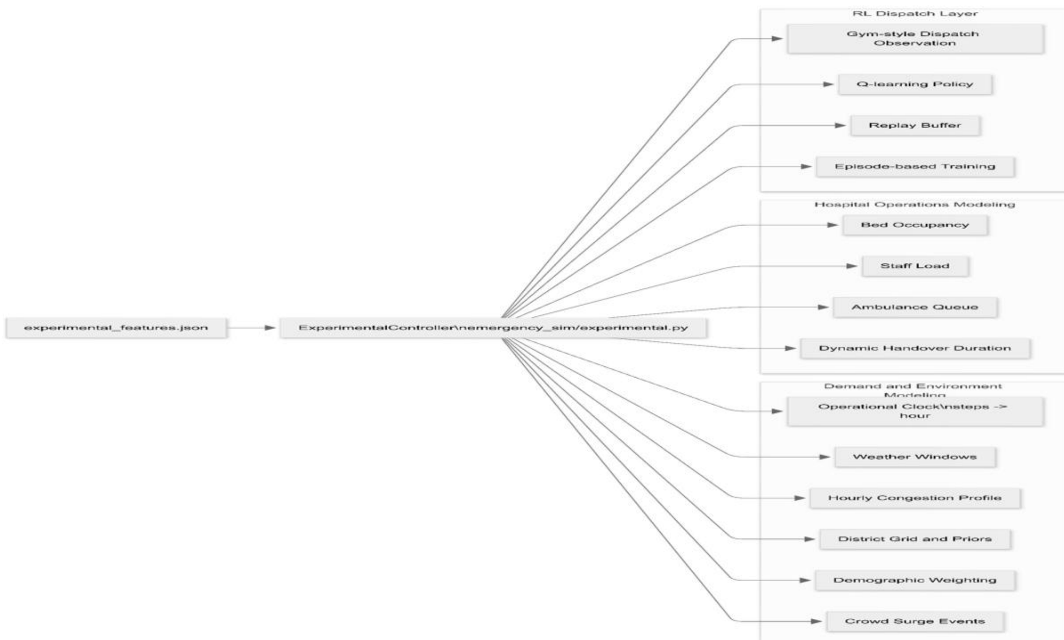

Fig. 5. Experimental Feature Layer

An optional reinforcement learning component is integrated into the dispatch system to improve the selection of primary ambulances. This learning module uses tabular Q-learning and works alongside the heuristic dispatch framework.

The formulation is defined as:

- State: representation of incident characteristics and candidate ambulance attributes (e.g., distance, availability, severity),
- Action: selection of a primary ambulance from a list of candidates,
- Reward: negative response time with additional penalties for delays or inefficient assignments.

The learning component does not replace the heuristic dispatcher; instead, it selects from the top-ranked candidates generated by the heuristic layer. This combined approach ensures stability while allowing incremental learning.

The learning process includes accumulating experience, replay-based updates, and optional checkpointing for persistence across simulation runs. This enables controlled experimentation with learning-based dispatch while maintaining baseline system reliability.

Metrics and Evaluation Pipeline
The system features a structured metrics pipeline for performance evaluation. During simulations, key operational data is recorded, including:

- Response time (from incident creation to arrival),
- incident completion rate,

- transport and handover durations,
- ambulance utilization and mission counts.

All metrics are exported in a structured JSON format at the end of each run. A post-processing dashboard analyzes these outputs, enabling visualization and comparison across different simulation configurations.

This pipeline supports reproducible experimentation by allowing controlled variations of parameters such as fleet size, incident frequency, and dispatch strategies. It also provides a quantitative basis for evaluating system performance and identifying trade-offs in emergency response operations.

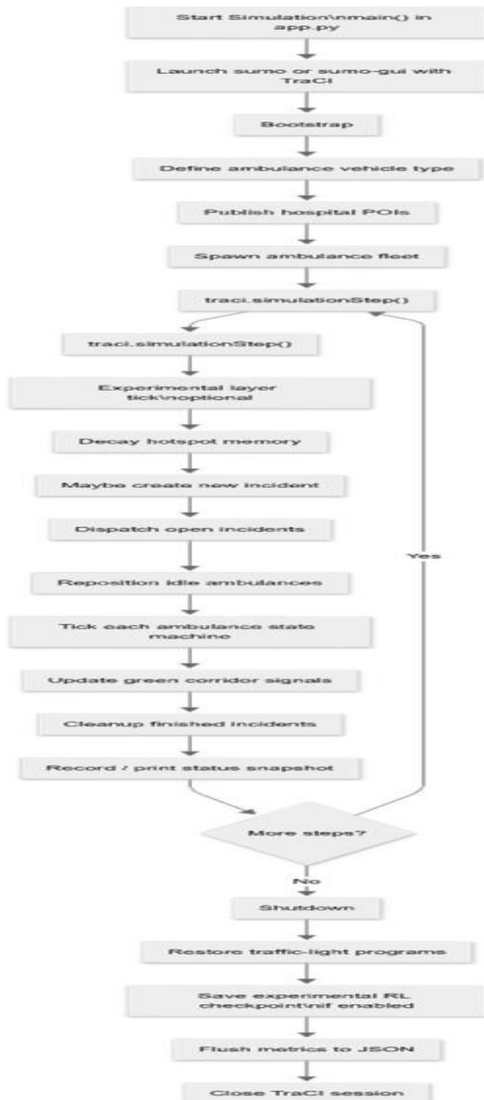

Fig. 6. *Runtime Control Loop*

Results and Discussion

This section presents the performance evaluation of the proposed hybrid multi-agent emergency response simulation platform under various configurations. The analysis focuses on key operational metrics, including response time, completion rate, and ambulance utilization, to assess the effectiveness of the dispatch and coordination mechanisms.

Experimental Setup

Experiments were conducted using the simulation environment built on SUMO (Simulation of Urban Mobility) with control via TraCI. The road network corresponds to an urban area in Pune.

The baseline configuration includes:

- Simulation duration: 1800 steps,

- ambulance fleet size: 4,
- incident interval: 90–180 steps,
- responders per incident: severity-dependent,
- green corridor mechanism: enabled.

Multiple runs were conducted with different random seeds to account for variability in traffic and incident generation.

Key Performance Metrics

The evaluation is based on the following metrics:

- Response Time: time from incident creation to first ambulance arrival,
- Completion Rate: ratio of completed incidents to total incidents,
- Transport Time: time from pickup to hospital arrival,
- Full Resolution Time: total time from incident creation to hospital handover,
- Ambulance Utilization: proportion of time ambulances are actively engaged in missions.

These metrics provide a thorough view of both efficiency and system load.

Baseline Performance

Under the default configuration, the system shows stable and consistent behavior. Most incidents receive timely responses, with high completion rates across runs. The heuristic dispatch mechanism effectively prioritizes high-severity and older incidents to ensure that critical cases are addressed quickly.

Response times are influenced by traffic conditions and the spatial distribution of ambulances. The hotspot-based patrol mechanism enhances coverage by redistributing idle ambulances toward high-demand areas, reducing average response delays over time.

Ambulance utilization stays balanced, with no single unit consistently overloaded. This indicates that the coverage-aware dispatch strategy successfully prevents resource clustering.

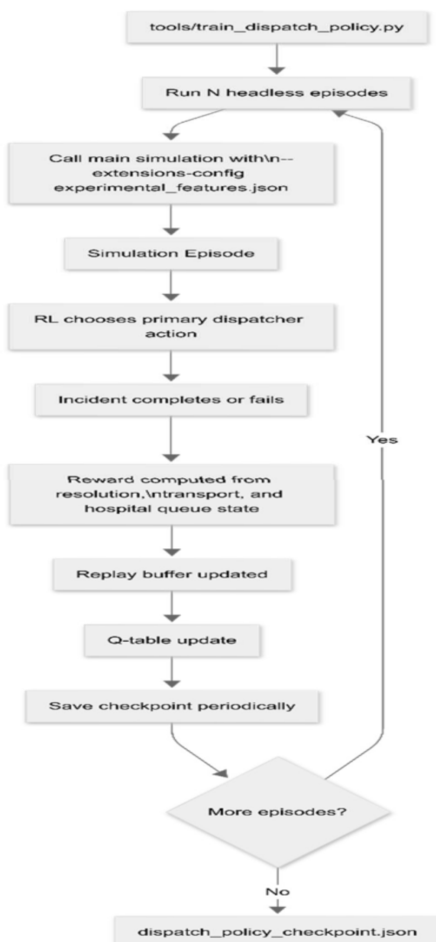

Fig. 7. *Training Workflow for the RL Dispatch Layer*

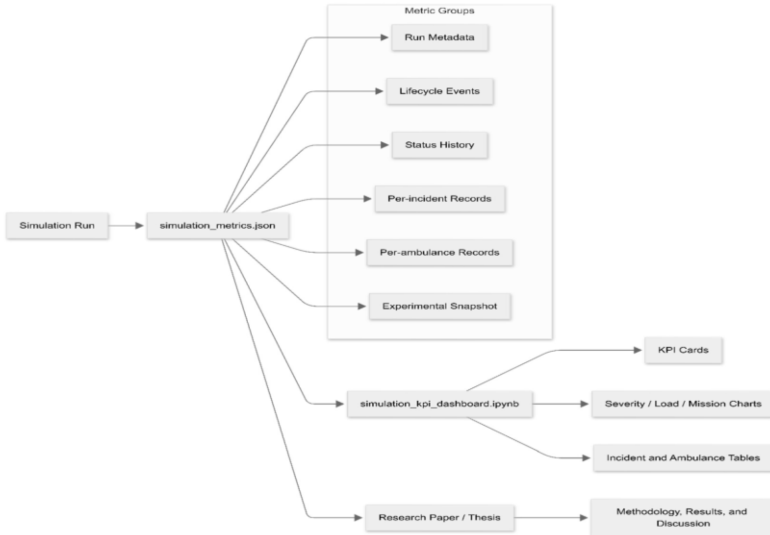

Fig. 8. *Research and Analytics Pipeline*

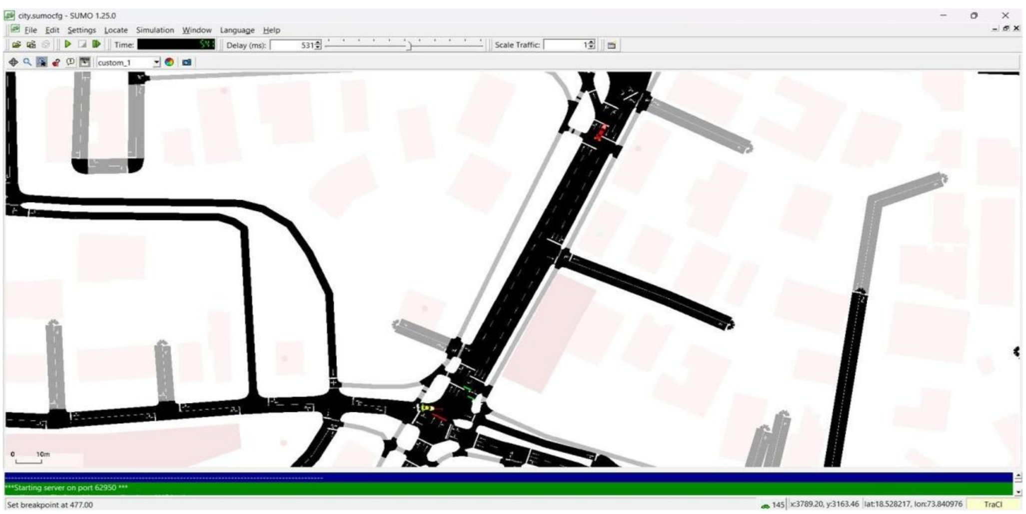

Fig. 9. *User Interface of Simulation Running*

Impact of Green Corridor Mechanism

Enabling traffic signal preemption leads to a significant reduction in travel time during active response and transport phases. Ambulances face fewer delays at intersections, resulting in quicker incident arrivals and shorter transport durations.

However, the impact is localized, as the current implementation only considers nearby traffic signals. While effective for short-range prioritization, the lack of corridor-level coordination limits its potential for overall optimization.

Dispatch Strategy Behaviour

The scoring-based dispatch model shows adaptive behavior in varying conditions. Incidents with higher severity and longer waiting times are consistently prioritized. Including a coverage penalty ensures that ambulances do not repeatedly draw from the same region, maintaining spatial balance.

Compared to a simple nearest-ambulance strategy, the heuristic approach performs better by taking into account multiple factors beyond distance. This leads to improved handling of simultaneous incidents and fewer service gaps.

Effect of Fleet Size

Increasing the number of ambulances improves response times and raises completion rates, as expected. With more resources available, the system handles concurrent incidents more efficiently and maintains better geographic coverage.

However, beyond a certain fleet size, the gains diminish. This suggests diminishing returns, indicating that optimal fleet sizing depends more on incident frequency and spatial distribution than simply increasing resources.

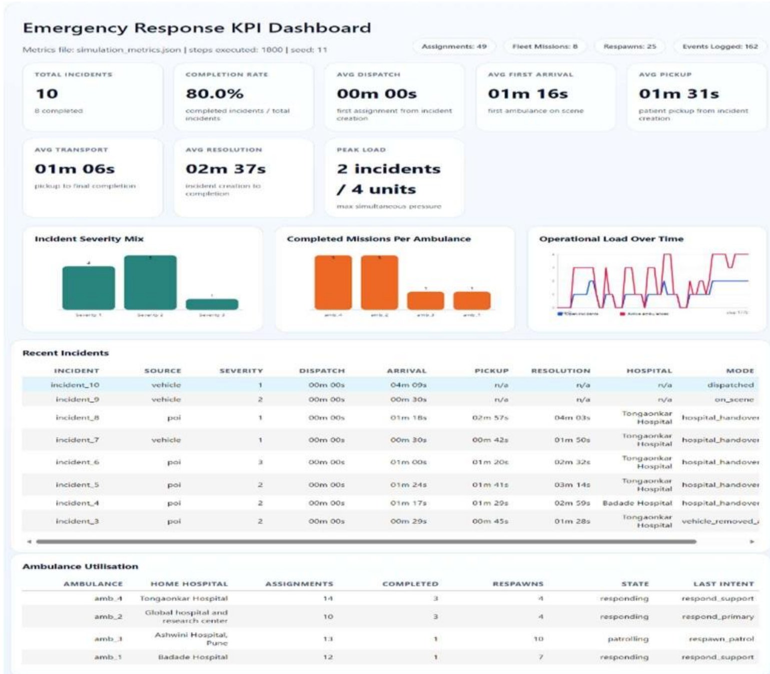

Fig. 10. *KPI Dashboard*

Observations on System Behavior

Several important behavioral patterns were observed:

- Hotspot Adaptation: Ambulances gradually concentrate in high-demand regions due to hotspot memory, improving responsiveness.
- Load Balancing: Coverage-aware dispatch prevents specific ambulances from being overused.
- Temporal Variability: Response performance changes with traffic density and incident timing.
- Robustness: The system operates stably even under multiple simultaneous incidents.

These observations highlight the effectiveness of combining centralized coordination with decentralized execution.

Discussion

The results show that the proposed hybrid architecture strikes a practical balance between efficiency and flexibility. The heuristic dispatch model, while not globally optimal, is computationally efficient and performs well under dynamic conditions.

The multi-agent behavior allows ambulances to adapt locally, while centralized dispatch ensures coordinated resource allocation. Adding traffic signal preemption further boosts system performance, especially in congested scenarios.

The optional learning component offers a pathway for future improvement, but its current impact is limited due to the use of tabular methods.

Overall, the platform serves as an effective testbed for evaluating emergency response strategies. It allows systematic experimentation and provides insights into the trade-offs between dispatch policies, resource allocation, and traffic conditions.

In summary, the system shows strong baseline performance, significant improvements through coordination mechanisms, and clear potential for further enhancement through advanced learning and optimization techniques.

Conclusion

This paper presented a hybrid multi-agent emergency response simulation platform for evaluating ambulance dispatch and coordination strategies in an urban environment. Built on a realistic road network and integrated with traffic simulation, the system combines centralized dispatch with decentralized agent autonomy to model the entire lifecycle of emergency response, from incident generation to hospital handover.

The results indicate that the proposed approach achieves stable and effective performance across key metrics such as response time, completion rate, and ambulance utilization. The scoring-based dispatch mechanism allows adaptive prioritization of incidents while maintaining spatial coverage, and the multi-agent design enables ambulances to respond flexibly to changing traffic conditions. Additionally, incorporating traffic signal preemption enhances travel efficiency during critical response phases.

A key contribution of this work is the development of a modular and extensible simulation framework that supports controlled experimentation. This platform allows systematic analysis of different operational parameters, including fleet size, dispatch strategies, and environmental conditions, making it suitable as a digital testbed for urban emergency response systems.

While the current implementation relies on heuristic decision-making with a limited reinforcement learning component, it lays a strong foundation for future research. Potential extensions include integrating deep reinforcement learning methods, optimizing corridor-level traffic signals, and calibrating using real-world emergency and traffic data.

In sum, this research proposes an integrated simulation platform for emergency response in urban contexts, bringing together simulation, multi-agent coordination, and applied artificial intelligence in a realistic, reproducible setting. We also emphasize the benefits of simulation-driven design as an appropriate first step towards implementing smart city solutions.

Future Scope

The proposed platform is a solid foundation on which further research can expand its realism, scalability, and value as an evaluation tool. Some future directions include:

1. Advanced Learning-Based Dispatch: The current system's dispatch policy is determined by a tabular reinforcement learning approach. This can be extended to use deeper reinforcement learning models, such as Deep Q-Networks or policy gradient approaches, which can support larger state spaces and more complex decisions. Multi-agent reinforcement learning could also be explored to co-optimize dispatch, patrol, and support unit actions.

2. Learning-Based Patrol and Rebalancing: The system currently employs a number of static and dynamic heuristic approaches, such as memory of hotspot areas, to govern the movement of ambulances. These can be replaced by learning-based policies capable of adapting to changing spatial and temporal demand patterns for optimal proactive deployment.

3. Corridor-Level Traffic Signal Optimization: The current green corridor feature focuses only on short-range effects of coordinated traffic signals. A potential future extension could incorporate a corridor-level approach where multiple intersections work together to produce larger, more reliable green waves for emergency vehicles in urban environments.

4. Real-World Data Integration: The current system's realism would be further enhanced by incorporating actual historical ambulance dispatch data, traffic flow, and hospital availability. This data can be used not only for a more accurate model, but also as a point of comparison and for the design of an actual decision support tool.

5. Enhanced Hospital Modeling

The existing hospital model can be extended to include:

- Department-level capacity (e.g., icu, emergency ward),
- Dynamic staff availability,
- Patient triage categories,
- And inter-hospital transfers.

Such enhancements would allow more realistic modeling of hospital congestion and its impact on emergency response decisions.

6. Multi-Objective Optimization: The proposed system's evaluation metrics could be extended to accommodate multi-objective optimization tasks that take into account a wider array of criteria beyond just response time, such as the quality of the served hospitals, geographic fairness, and operational cost.

7. Scalability and Larger Urban Networks: The current simulation focuses on a small-scale network in New York City. Extension of the system to a full-fledged urban region or a number of interconnected regions would pose challenges related to computational complexity that are worth studying. Techniques for parallel and distributed simulation should be considered to address this challenge.

8. Integration with IoT and Real-Time Systems: In addition to real-time traffic information, the system can be augmented to connect to location-based sensors and smart devices via an IoT framework. The information collected from these devices can be used to improve real-time rerouting decisions of ambulances and guide decisions at the decision-support level.

9. Human-in-the-Loop Decision Support: Finally, the system can be augmented to support human dispatchers in their decision-making role. A user interface for interacting with the system and directly providing input to the dispatchers' decision processes would not only be valuable for evaluation and training but would bring it closer to practical use in a dispatch center environment.

The aforementioned research directions can extend the current simulation into a fully capable smart city decision support system for emergency response.

References

1. SUMO (Simulation of Urban Mobility), "Simulation of Urban Mobility (SUMO)," Available: https://www.eclipse.org/sumo/
2. TraCI, "TraCI: Traffic Control Interface," Available: https://sumo.dlr.de/docs/TraCI.html
3. OpenStreetMap Contributors, "OpenStreetMap," Available: https://www.openstreetmap.org/
4. Wooldridge, M., *An Introduction to Multi-Agent Systems*, 2nd ed., Wiley, 2009.
5. Sutton, R. S., & Barto, A. G., *Reinforcement Learning: An Introduction*, 2nd ed., MIT Press, 2018.
6. Gendreau, M., Laporte, G., & Semet, F., "The Maximal Covering Location Problem for Emergency Vehicles," *European Journal of Operational Research*, vol. 147, no. 3, pp. 499–518, 2003.
7. Schmid, V., "Solving the Dynamic Ambulance Relocation and Dispatching Problem Using Approximate Dynamic Programming," *European Journal of Operational Research*, vol. 219, no. 3, pp. 611–621, 2012.
8. Chen, X., He, F., Yin, Y., & Du, Y., "Optimal Design of Traffic Signal Control for Emergency Vehicle Preemption," *Transportation Research Part B*, vol. 99, pp. 134–150, 2017.
9. Krajzewicz, D., Erdmann, J., Behrisch, M., & Bieker, L., "Recent Development and Applications of SUMO – Simulation of Urban Mobility," *International Journal On Advances in Systems and Measurements*, 2012.
10. Maciejewski, M., & Nagel, K., "Towards Multi-Agent Simulation of the Dynamic Vehicle Routing Problem in MATSim," *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2013.
11. Zhang, R., Pavone, M., & Smith, S. L., "Control of Robotic Mobility-on-Demand Systems: A Queueing-Theoretical Perspective," *International Journal of Robotics Research*, 2016.
12. Van der Pol, E., Oliehoek, F. A., & Whiteson, S., "Coordinated Deep Reinforcement Learners for Traffic Light Control," *Proceedings of NIPS Workshop*, 2016.
13. Abouee-Mehrizi, H., Baron, O., & Berman, O., "Ambulance Location and Relocation Models with Time-Dependent Travel Times," *European Journal of Operational Research*, 2017.
14. Li, Y., Zheng, Y., Zhang, H., & Chen, L., "Traffic Prediction in a Bike-Sharing System," *Proceedings of SIGSPATIAL*, 2015.