Proceedings of the International Conference on Industrial Engineering and Operations Management

Publisher: IEOM Society International, USA DOI: 10.46254/AP07.20260224
Published: March 25, 2026

Optimizing Ambulance Routing with Time Series Forecasting for Hospital Readiness and Emergency Call Patterns

Asif Salman Khan Tanim, Zahin Ar Rafi and Ferdous Sarwar
Department of Industrial and Production Engineering
Bangladesh University of Engineering and Technology
Dhaka, Bangladesh
2008002@ipe.buet.ac.bd , zahin@ipe.buet.ac.bd , ferdoussarwar@ipe.buet.ac.bd

Abstract

Effective ambulance routing is a top priority when it comes to getting emergency help out to people in need. The traditional go-to approach when it comes to ambulance routing has been to just focus on getting the ambulance to the scene in the shortest time possible, without much thought given to whether the hospital it's headed for is already overwhelmed with patients or doesn't have the right resources that need critical care. This study presents a novel, data-driven approach that optimizes ambulance routing by integrating real-time hospital crowdedness forecasting with intelligent routing strategies. Rather than focusing solely on minimizing travel time, this framework takes into account hospital overcrowding, bed availability, and specialized care needs. A hybrid forecasting model is employed to predict hospital crowdedness levels, combining ARIMA (for capturing trends and seasonality), LSTM neural networks (for addressing complex patterns), and Gradient Boosting (for robustness). A Deep Q-Network (DQN) agent, trained over 500 episodes, learns to make dynamic routing decisions, balancing travel distance with hospital readiness. The agent operates within a Markov Decision Process (MDP) to guide ambulances to the hospitals best equipped for specific emergencies. The model will be tested in Montgomery County, PA, using real-world data to evaluate its effectiveness. By optimizing both ambulance routing and hospital preparedness, this innovative approach aims to significantly improve emergency care delivery, ensuring faster, smarter, and more efficient treatment for patients in critical need.

Keywords
Ambulance Routing, Time Series Forecasting, Hospital Crowdedness, Deep Q-Network, ARIMA

1. Introduction
Every year, millions of people call emergency services expecting fast, effective help. What they don't expect and what they rarely know is that the ambulance carrying them may be headed to a hospital already overwhelmed with patients. When that happens, the speed of the ambulance means little. The bottleneck isn't on the road. It's at the ER door.

Time-sensitive emergencies like cardiac arrest, traumatic injury, respiratory failure account for nearly 15% of global deaths and disability-adjusted life years (Vos et al., 2020). Emergency Medical Services (EMS) are the critical bridge between the moment of crisis and the moment of care. Response time, how quickly an ambulance reaches the patient, is the metric most studied and most optimized (Srinath, 2018). But response time is only half the equation. The other half, largely ignored by existing systems, is what happens when the ambulance arrives at a hospital that isn't ready.
The goal of all previous ambulance routing research has been to minimize travel time while ignoring the capacity and crowding of the destination hospital. Patients may arrive at overcrowded emergency rooms, experience longer wait times, and ultimately have worse clinical outcomes if an ambulance is routed to the geographically closest hospital in

high-volume urban environments without taking current or anticipated occupancy into account. A major and unexplored domain in the literature on EMS optimization is this capacity-blindness.

The current study addressed this gap and simultaneously takes accident location and hospital crowdedness index into account when figuring out the best ambulance route. The main goal is to develop an advanced ambulance routing system that improves patient care by selecting hospitals based on real-time crowding and hospital location. In order to determine the best route schedule for the ambulance, this study takes into account both the location of the accident and a predictive hospital crowdedness index that is produced by a hybrid ARIMA-LSTM forecasting pipeline. A more efficient ambulance routing system that dynamically adjusts routes based on emergency call data and anticipated hospital crowding is the expected result. This allows for proactive routing to less overburdened hospitals before they become overloaded.

This study makes four distinct contributions. In order to create a hybrid forecasting pipeline that is better than either approach alone, a novel hospital crowdedness forecasting model is first proposed. This model combines the interpretable trend-and-seasonality decomposition of ARIMA with the ability of LSTM to model non-linear residual patterns. Second, a carefully crafted reward function that reflects the multi-objective nature of actual EMS dispatching is used to formally model the routing problem as a Markov Decision Process. Third, a Deep Q-Network agent exhibits better load-balancing characteristics than both the nearest-hospital heuristic and a traditional weighted scoring baseline after being trained and verified against real historical call data. Fourth, the entire pipeline is put through five simulated emergency scenarios to test its resilience to hospital overloads and mass casualty incidents.

2. Problem Formulation and Literature Review

The problem addressed in this work is formally stated as follows: given a set of N hospitals with known geographic coordinates, bed capacities, trauma capabilities, and dynamically varying crowdedness indices, and given a continuous stream of georeferenced 911 emergency calls each characterized by location, time, and emergency type, determine, for each incoming call, the optimal hospital to which the dispatched ambulance should be routed, such that the expected response time is minimized, the probability of routing to an overcrowded hospital is minimized, patient-hospital type compatibility (i.e., trauma patients to trauma-capable facilities) is maximized, and the overall load is distributed equitably across all available hospitals.

This is a multi-objective, time-varying, stochastic optimization problem. The stochasticity arises from the uncertain future crowdedness state of each hospital, the unpredictable spatial and temporal distribution of emergency calls, and the dynamic feedback between routing decisions and hospital occupancy. The time-varying nature arises because crowdedness evolves continuously as patients are admitted and discharged, and because emergency call volume exhibits strong temporal patterns (time of day, day of week, seasonal trends).

A comprehensive review of prior work is organized by methodological paradigm below.

2.1 Classical Optimization Approaches

The earliest formulations of the EMS resource allocation problem drew on operations research. The Maximal Covering Location Problem (MCLP) (Jayaraman, 1994) maximizes the population covered within a specified response time threshold, treating hospital locations as fixed. The double-standard model (Gendreau et al., 1997) was proposed to handle cases where a patient can be reached by multiple vehicles, adding redundancy constraints. These static approaches do not account for time-varying demand or hospital capacity. Simulation-based optimization frameworks have since been proposed to address the dynamic nature of ambulance deployment and relocation (Zhen et al., 2014). The equitable distribution of workload across EMS stations has also been studied through load-balancing location models that constrain the population allocated per station (Jánošíková et al., 2015).

Graph-based routing algorithms were subsequently applied to EMS. Arunmozhi and William (2014) proposed the use of Dijkstra's shortest-path algorithm for real-time ambulance routing. Michael and Xavier (2018) applied the A* algorithm, demonstrating improved efficiency in dynamic road networks. Both approaches focus exclusively on minimizing travel distance or time, with no consideration of destination hospital state.

Kamireddy and Keshavamurthy (2016) introduced a clustering-based approach using K-means and density-based clustering to segment emergency zones and assign ambulances to the closest relevant hospital. While this improves spatial efficiency, the assignment remains static and does not adapt to hospital-level crowdedness. Tavakkoli-Moghaddam et al. (2018) proposed a bi-objective location-allocation model for disaster scenarios, incorporating ambulance routing alongside facility placement. However, the model targets disaster response planning rather than continuous operational dispatch, and does not integrate real-time hospital occupancy data.

2.2 Time Series Forecasting in Healthcare Operations

Forecasting hospital demand and occupancy has a rich literature in healthcare management. ARIMA models have been widely applied to emergency department (ED) visit prediction (Setzler et al. 2009). Autoregressive Integrated Moving Average models are well-suited to stationary or trend-stationary time series and have demonstrated reliable short-horizon forecasting in multiple hospital systems. The selection of ARIMA order parameters (p, d, q) via the Akaike Information Criterion (AIC) is a standard and well-validated practice (Box et al., 1970).

Recurrent neural networks, and particularly Long Short-Term Memory (LSTM) architectures introduced by Hochreiter and Schmidhuber (1997), have demonstrated superior performance over ARIMA on non-linear time series with long-range dependencies. LSTM networks use gating mechanisms (input, forget, and output gates) to selectively retain or discard information across time steps, overcoming the vanishing gradient problem that limits simple RNNs. In healthcare operations, LSTM models have been applied to ICU occupancy prediction, patient length-of-stay modelling, and emergency call volume forecasting (Yu et al. 2021).

Hybrid ARIMA-LSTM models, in which ARIMA captures the linear component of a time series and LSTM corrects the non-linear residuals, have been shown to outperform either model individually across a range of forecasting tasks (Zhang 2003). This two-stage decomposition approach is adopted in the present work. When LSTM is unavailable, Gradient Boosting Regression on lag features of ARIMA residuals serves as a theoretically grounded and practically robust alternative.

2.3 Reinforcement Learning for Routing Optimization

Reinforcement learning (RL) has emerged as a powerful paradigm for sequential decision-making under uncertainty. Mnih et al. (2015) introduced Deep Q-Networks (DQN), demonstrating human-level performance across Atari games by combining Q-learning with deep neural function approximation. DQN employs two stabilizing techniques: experience replay by storing transitions in a buffer and sampling mini-batches to break temporal correlations and a target network that is updated periodically to provide stable training targets.

The application of RL to ambulance dispatching and routing is an active research frontier. Several recent works have modelled EMS dispatch as a Markov Decision Process, with states representing vehicle location and demand patterns, and rewards based on response time reduction (Liu et al. 2020; Ji et al. 2019). However, no prior work has simultaneously integrated hospital crowdedness forecasting into the RL state representation, enabling the agent to reason about both travel time and destination hospital state within a unified decision-making framework. This represents the primary research gap addressed by the present study.

2.4 Research Gap and Contribution

The literature review reveals a clear bifurcation: routing optimization methods (classical or RL-based) treat hospital state as static or ignore it entirely, while forecasting methods predict hospital demand without coupling their predictions to a routing decision layer. The present work bridges this gap by designing a complete pipeline in which ARIMA-LSTM crowdedness forecasts are embedded into the RL state space, creating a dynamic, anticipatory routing system. The classical weighted scoring baseline further enables ablation analysis, isolating the contribution of the learned RL policy from the forecasting component itself.

3. Methods

The proposed framework consists of four tightly integrated components: (i) a data preprocessing and feature engineering pipeline, (ii) a hybrid hospital crowdedness forecasting model, (iii) a Deep Q-Network routing agent operating within a Markov Decision Process environment, and (iv) a real-time simulation system that couples forecasting outputs with routing decisions and implements an adaptive feedback loop (Table 1).

Table 1. Summary of Overall Architecture

Stage	Component
Stage 1	Data Preprocessing and Feature Engineering
Stage 2	Hospital Crowdedness Forecasting (Hybrid Arima + LSTM)
Stage 3	Ambulance Touting via Deep Q-Network
Stage 4	Real-Time Simulation with Adaptive Feedback

3.1 Data Preprocessing and Feature Engineering

Four primary datasets were ingested: a hospital crowdedness index aggregated by hospital, month, day-of-week, and time period; a hospital static characteristics dataset containing bed counts, ratings, trauma designation, and geographic coordinates; a raw 911 call dataset with timestamps, GPS coordinates, and call reasons; and an enriched routing dataset that augments each call with nearest-hospital assignment, crowdedness at time of call, a recommended hospital produced by a baseline routing system, and a diversion flag.

Data cleaning proceeded in three steps. First, numeric hospital attributes were imputed using the column median (a strategy robust to the skewed distributions typical of healthcare data). Second, duplicate records in the crowdedness index were identified and removed. Third, emergency call records with missing GPS coordinates were excluded as they cannot be meaningfully routed. Extreme crowdedness values were filtered using the 3×IQR rule: observations more than three interquartile ranges below the first quartile or above the third quartile were treated as data artifacts and removed.

Feature engineering introduced six new variables. The IsHoliday binary flag marks the eight major U.S. public holidays within the dataset window (January 2016 to December 2017), as holiday periods are associated with elevated emergency call volumes and reduced hospital staffing. The RushHour flag marks morning (07:00–09:59) and evening (16:00–18:59) rush periods, capturing the dual effect of increased road accidents and emergency service demand. Emergency calls were categorized into five types (Cardiac, Trauma, Respiratory, Neurological, and Other), enabling the routing agent to match patient acuity to hospital capability. For the crowdedness time series, three temporal smoothing features were derived: rolling 3-step mean (CrowdRoll3), rolling 7-step mean (CrowdRoll7), and an exponentially weighted moving average with a span of 7 steps (CrowdEWM). Two lag features (CrowdLag1 and CrowdLag3) capture direct autoregressive dependencies at one and three steps, respectively.

3.2 Hospital Crowdedness Forecasting

The normalized crowdedness index for each hospital at a given time step is defined as the ratio of concurrent call count to hospital bed capacity, normalized to the unit interval [0, 1] using global min-max scaling. The forecasting task is to predict this index over a seven-step horizon for each hospital independently.

Step 1: ARIMA Forecasting

The Augmented Dickey-Fuller (ADF) test is applied to verify the stationarity of each hospital's crowdedness series. Autocorrelation (ACF) and Partial Autocorrelation (PACF) plots guide the selection of tentative (p, q) orders. The pmdarima auto_arima function is then used to perform a stepwise grid search over $p \in \{0,...,4\}$, $q \in \{0,...,4\}$, $d \in \{0,1\}$, selecting the model with the minimum AIC. Model training uses 80% of the chronologically ordered time series, and evaluation proceeds via a rolling one-step-ahead forecast on the remaining 20% — the most operationally faithful evaluation strategy, as it mirrors deployment conditions where only past observations are available.

$ARIMA(p,d,q)$: $(1 - \sum_{i=1}^{p} \varphi_i L^i)(1-L)^d X_t = (1 + \sum_{j=1}^{g} \theta_j L^j)\varepsilon_t$

where L is the lag operator, φ_i are autoregressive coefficients, θ_j are moving-average coefficients, d is the degree of differencing, and $\varepsilon_t \sim N(0,\sigma^2)$ is white noise.

Step 2: LSTM Forecasting

When TensorFlow is available, a two-layer stacked LSTM network is constructed with a sliding window of W = 7 time steps as input. The first LSTM layer has 64 units with return_sequences=True, followed by a Dropout layer (rate 0.2). The second LSTM layer has 32 units with return_sequences=False, followed by another Dropout layer (rate 0.2). A Dense layer with 16 ReLU units and a final scalar Dense output constitute the prediction head. The model is compiled with the Adam optimizer (learning rate 1×10^{-3}) and mean squared error loss. EarlyStopping with patience of 10 epochs (monitoring validation loss) and ReduceLROnPlateau with a factor of 0.5 prevent overfitting and facilitate convergence.

LSTM cell: $f_t = \sigma(Wf\cdot[h_{t-1},x_t] + bf)$; $i_t = \sigma(Wi\cdot[h_{t-1},x_t] + bi)$

$\tilde{c}_t = \tanh(Wc\cdot[h_{t-1},x_t]+bc)$; $c_t = f_t\odot c_{t-1} + i_t\odot\tilde{c}_t$; $h_t = o_t\odot\tanh(c_t)$

Step 3: Hybrid Model

The hybrid model decomposes the crowdedness time series into a linear component (captured by ARIMA) and a non-linear residual component (captured by LSTM or, as a fallback, Gradient Boosting Regression). ARIMA is first fit to the full training series to obtain in-sample fitted values. Residuals are computed as the difference between observed and ARIMA-fitted values. The residual series is then scaled to [-1, 1] and used to train a secondary LSTM or GBR model. At inference time, the hybrid prediction is the sum of the ARIMA forecast and the residual model's output.

$\hat{y}_t = ARIMA\_forecast(t) + Residual\_model\_forecast(t)$

Model Performance Assessment: Model performance will be assessed using four evaluation metrics, reported separately for each hospital. These are defined in Table 2 below.

Table 2. Performance Evaluation Metrics for the Crowdedness Forecasting Model

Metric	Full Name	Interpretation
R^2	Coefficient of determination	Proportion of variance explained; higher is better (→ 1)
MAE	Mean Absolute Error	Average absolute prediction error; same unit as target
RMSE	Root Mean Squared Error	Penalizes larger errors more heavily than MAE
MAPE	Mean Absolute Percentage Error	Scale-independent error; useful for cross-stream comparison

3.3 Ambulance Routing as a Markov Decision Process

The routing problem is cast as a finite-horizon Markov Decision Process (MDP) defined by the tuple (S, A, P, R, γ), where S is the state space, A is the action space, P is the transition probability function, R is the reward function, and γ is the discount factor.

State Space: The state vector $s_t \in \mathbb{R}^{(4N+2)}$ for an incoming call is constructed as:

$s_t = [lat/90, lon/180, c_1,...,c_n, b_1,...,b_n, \tau_1,...,\tau_n, d_1/100,...,d_n/100]$

Where lat, lon are the normalized GPS coordinates of the emergency call origin; $c_i \in [0,1]$ is the current crowdedness estimate for hospital i (from the ARIMA forecast); $b_i \in [0,1]$ is the normalized bed count (beds / max beds across all hospitals); $\tau_i \in \{0,1\}$ indicates trauma center designation; and d_i is the geodetic distance in km to hospital i, normalized by dividing by 100. For N = 7 hospitals, the state vector has dimensionality 4(7)+2 = 30 (Table 3).

Table 3. State Space Components of the MDP Routing Agent

Component	Description	Normalization	Dimensionality
lat, lon	Call GPS coordinates	Normalized: lat/90, lon/180	2 dimensions
$c_1 \dots c_n$	Hospital crowdedness	ARIMA/hybrid forecast ∈ [0, 1]	N dimensions
$b_1 \dots b_n$	Normalized bed count	Beds / max beds ∈ [0, 1]	N dimensions
$\tau_1 \dots \tau_n$	Trauma designation	Binary ∈ {0, 1}	N dimensions
$d_1 \dots d_n$	Distance to hospital	Geodetic km / 100	N dimensions
Total state vector (N=7 hospitals)	4(7) + 2 = 30 dimensions		

Action Space: The action $a_t \in \{0, 1, ..., N-1\}$ is an integer index identifying the selected hospital. This discrete action space is well-suited to DQN.

Reward Function: The reward r_t is a weighted linear combination of four factors:

$$r_t = -0.40\cdot(d_a/50) - 0.40\cdot c_a + 0.15\cdot b_a + 0.05\cdot(\tau_a \cdot \tau\_need)$$

where d_a is the distance in km to the chosen hospital a, c_a is its current crowdedness, b_a is its normalized bed count, τ_a is its trauma designation, and $\tau\_need \in \{0,1\}$ indicates whether the call requires trauma capability. The weights (0.40, 0.40, 0.15, 0.05) reflect the relative importance of distance and crowdedness avoidance versus capacity and type matching, and were set based on clinical priority literature for EMS dispatch (Table 4).

Table 4. Components of the MDP Reward Function

Factor	Rationale	Weight	Term
Distance penalty	Penalizes routing to farther hospitals	-0.40	$-0.40 \times (d_a / 50)$
Crowdedness penalty	Penalizes routing to busier hospitals	-0.40	$-0.40 \times c_a$
Capacity bonus	Rewards hospitals with higher bed counts	+0.15	$+0.15 \times b_a$
Trauma match bonus	Rewards trauma-capable routing when needed	+0.05	$+0.05 \times (\tau_a \cdot \tau\_need)$

3.4 Deep Q-Network Agent

The DQN agent maintains two neural networks: a Q-network $Q\theta(s,a)$ with parameters θ that is updated at each training step, and a target network $Q\bar{\theta}(s,a)$ with parameters $\bar{\theta}$ that are periodically synchronized with θ (every 50 steps). Both networks are 4-layer feedforward networks with architecture [state_dim → 128 → 64 → 32 → action_dim], ReLU activations on hidden layers, and a linear output layer (Table 5).

Table 5. Deep Q-Network Architecture and Hyperparameters

Component	Specification	Detail
Q-network & Target	4-layer feedforward	[state_dim → 128 → 64 → 32 → action_dim]; ReLU hidden; linear output
Experience Replay	Circular buffer, size 10,000	Mini-batch of 64 transitions sampled uniformly
Target Sync Interval	Every 50 steps	Parameters $\bar{\theta} \leftarrow \theta$; stabilizes TD targets
Discount Factor	$\gamma = 0.95$	Balances immediate vs future reward
Learning Rate	$\alpha = 1\times10^{-3}$	Adam optimizer
Exploration	$\varepsilon = 1.0 \rightarrow 0.05$	Multiplicative decay 0.995/episode; ≈ 598 episodes to reach floor

Experience Replay: Transitions (s_t, a_t, r_t, s_{t+1}, $done_t$) are stored in a circular buffer of capacity 10,000. At each learning step, a mini-batch of 64 transitions is sampled uniformly at random. This decorrelates sequential experiences and stabilizes gradient estimates.
Temporal Difference Update: For each transition in the mini-batch, the TD target is computed as:

$$y_i = r_i + \gamma \cdot max\_a' \; Q\bar{\theta}(s_{i+1}, a') \quad (if\ not\ terminal)$$

$$y_i = r_i \; (if\ terminal)$$

The Q-network parameters are updated by minimizing the mean squared TD error: $L(\theta) = E[(y_i - Q\theta(s_i, a_i))^2]$. The discount factor $\gamma = 0.95$ and learning rate $\alpha = 1\times10^{-3}$ are used throughout.
Epsilon-Greedy Exploration: The agent begins with $\varepsilon = 1.0$ (fully random exploration) and decays ε multiplicatively at rate 0.995 per episode, reaching a floor of $\varepsilon\_min = 0.05$ after approximately 598 episodes. This ensures sufficient early exploration while converging to exploitation of learned Q-values.

3.5 Classical Routing Baseline

For comparison, a deterministic weighted scoring baseline is implemented. For each incoming call, a composite score is computed for every hospital:

$$score(h) = -w\_d\cdot(d\_h/50) - w\_c\cdot c\_h + w\_b\cdot(beds\_h/max\_beds) + w\_t\cdot(\tau\_h \cdot \tau\_need)$$

with fixed weights w_d = 0.40, w_c = 0.40, w_b = 0.15, w_t = 0.05. The hospital with the highest score is selected. This baseline uses the same multi-objective formulation as the DQN reward function but applies fixed weights without learning from experience.

3.6 Real-Time Simulation with Adaptive Feedback

The RealTimeRoutingSystem class simulates operational deployment. For each call, the DQN agent (or classical baseline) selects a hospital. After each routing decision, the system slightly increases the crowdedness estimate for the chosen hospital by a factor $\alpha = 0.05$, clamped to [0, 1]. This adaptive feedback loop models the actual pressure that routed patients exert on hospital capacity, allowing the simulation to capture the emergent load-balancing dynamics of the full pipeline. A diversion flag is raised and logged whenever a routed hospital's crowdedness exceeds the threshold of 0.75.

3.7 Integration of Forecasting and Routing Components

The sequential integration of the hybrid forecasting model with the DQN routing agent reduces the uncertainty inherent in static crowdedness estimation that challenges conventional rule-based dispatch systems. Forecast-updated crowdedness values replace static historical averages in the state vector, enabling the DQN to respond dynamically to anticipated hospital load conditions over the 7-step forward horizon. This closed-loop, data-driven approach supports the overarching objective of minimizing both patient transport time and hospital overcrowding simultaneously, while the adaptive feedback simulation validates the real-world feasibility of the learned routing policy.

4. Data Collection

This study is based on four complementary datasets drawn from real-world emergency medical services operations in Montgomery County, Pennsylvania, United States, covering the period from January 2016 to December 2017.

4.1 Dataset Descriptions

Dataset 1: Hospital Static Characteristics (hospitals_clean.csv). This dataset contains one record per hospital (N = 7) with the following attributes: hospital name, city, overall CMS quality rating (1–5), mortality and readmission national comparison scores, patient experience score, geographic coordinates (latitude, longitude), number of licensed beds, trauma center designation and category, and a binary HasTraumaCenter indicator. These attributes populate the static components of the MDP state vector and the classical scoring function.

Dataset 2: 911 Call Records (calls_clean.csv). This dataset contains 30,079 individual emergency call records, each with GPS coordinates, township, date, time, and the free-text reason for calling. After preprocessing, temporal features including Year, Month, Day, Hour, DayOfWeek, IsWeekend, TimePeriod, IsHoliday, RushHour, and Quarter were derived. Call reasons were mapped to five clinical categories (Cardiac, Trauma, Respiratory, Neurological, and Other) using keyword matching on the free-text field.

Dataset 3 (Combined and Modified from Dataset 1and 2):

Hospital Crowdedness Index (hospital_crowdedness_index.csv). This dataset contains 1,371 records aggregated at the level of hospital × month × day-of-week × time-period. Each record reports the count of emergency calls directed to that hospital during that time period (CallCount), as well as static hospital characteristics (HospitalBeds, HospitalRating, HasTrauma) and two crowdedness metrics: the raw CrowdednessIndex (defined as CallCount / HospitalBeds, representing occupancy pressure) and CrowdednessIndex_Norm (linearly scaled to [0,1] for model input). Time periods are defined as four daily blocks: Morning (00:00–05:59), Afternoon (06:00–11:59), Evening (12:00–17:59), and Night (18:00–23:59). This dataset serves as the primary target series for the forecasting models.
Dataset 4(Combined and Modified from Dataset 1,2 and 3):

Enriched Routing Dataset (calls_enriched_with_routing.csv). This is the primary evaluation dataset, containing all 30,079 call records augmented with: NearestHospital (the nearest by geodetic distance, representing the traditional baseline), DistanceToHosp_km (geodetic distance to the nearest hospital), hospital attributes at time of call (HospitalBeds, HospitalRating, HospitalHasTrauma, CrowdednessIndex, CrowdednessIndex_Norm), a RecommendedHospital (derived from an optimization-aware routing algorithm applied retrospectively), RoutingScore (the composite score of the recommended hospital), RecommendedDist_km (distance to the recommended hospital), and DiversionNeeded (a binary flag set to 1 when the nearest hospital's crowdedness exceeded the diversion threshold at time of call). This dataset enables direct comparison between traditional nearest-hospital routing and the optimized recommendation.

4.2 Dataset Statistics

Across 30,079 emergency calls, the five most common call reasons are falls/trauma-related incidents, cardiac and chest pain events, respiratory distress, abdominal pain, and altered mental status, a distribution consistent with national EMS utilization patterns. The temporal distribution shows clear weekday-peaking demand patterns, with Fridays recording the highest call volume and Sunday afternoons the lowest. Rush-hour effects are pronounced, with calls between 08:00–09:00 and 17:00–18:00 elevated by approximately 18% relative to the daily mean. The diversion rate in the enriched dataset, the fraction of calls for which the recommended hospital differs from the nearest hospital is 62.3%, underscoring the practical significance of capacity-aware routing.

The seven hospitals range in size from 126 to 665 licensed beds, with a mean of 337.7 beds. Two hospitals have designated trauma centers. Hospital overall CMS ratings range from 1 to 4. The mean normalized crowdedness index across all hospitals and time periods is 0.218, with a standard deviation of 0.095, indicating that while average crowdedness is moderate, peak crowdedness exceeds 0.70 in 12.4% of all observed time periods, a clinically meaningful proportion that motivates proactive routing.

5. Results and Discussion

The results tell a clear story: a smarter routing system saves time, reduces pressure on hospitals, and does so without asking ambulances to travel unreasonably far. This section unpacks that story in four parts consisting forecasting accuracy, routing performance, stress-scenario resilience, and economic impact.

The single most important number from Table 8 is this: under traditional nearest-hospital routing, 62.3% of calls are sent to an already-pressured hospital. More than three in five patients are routed somewhere that may not be ready for them. The DQN agent cuts that figure to 15.2%, a 47-percentage-point reduction, by learning to look ahead at hospital conditions rather than just minimizing distance.

5.1 Numerical Results

Table 6 presents the forecasting accuracy of the three crowdedness prediction models: ARIMA, LSTM, and the Hybrid (ARIMA + LSTM / Gradient Boosting), evaluated on the held-out 20% test set of the busiest hospital's crowdedness index series. All metrics are computed in the original [0,1] normalized scale.

Table 6. Forecasting model performance

Model	RMSE	MAE	MAPE (%)
ARIMA	0.0412	0.0318	14.82
LSTM (Keras)	0.0387	0.0291	13.45
Hybrid (ARIMA + LSTM)	0.0341	0.0263	12.17
Hybrid (ARIMA + GBR Fallback)	0.0358	0.0278	12.91

The Hybrid model achieves the lowest error across all three metrics, with RMSE = 0.0341, MAE = 0.0263, and MAPE = 12.17%, representing improvements of 17.2%, 17.3%, and 17.9% respectively over the standalone ARIMA baseline. The LSTM model also outperforms ARIMA when TensorFlow is available (RMSE improvement of 6.1%), confirming the presence of non-linear components in the crowdedness series that ARIMA cannot capture. The Gradient Boosting fallback hybrid achieves performance intermediate between LSTM-hybrid and standalone ARIMA, validating the robustness of the two-stage residual correction philosophy even without deep learning infrastructure.

The MAPE values, ranging from 12.17% to 14.82%, are consistent with the ranges reported in the hospital demand forecasting literature (typically 10%–20%), and are sufficiently accurate for the routing application, where forecast errors of less than ±0.05 crowdedness units have negligible impact on routing decisions.

Table 7. ARIMA grid search results: top 5 configurations sorted by RMSE.

Metric	ARIMA Order	AIC	RMSE (Test)
Best configuration	(1,1,1)	-124.3	0.0412
2nd best	(2,1,1)	-122.7	0.0419
3rd best	(1,1,2)	-121.9	0.0428
4th best	(0,1,1)	-119.4	0.0451
5th best	(2,1,0)	-118.2	0.0463

Table 7 presents the results of the ARIMA grid search across all combinations of $p \in \{0,1,2\}$, $d \in \{0,1\}$, $q \in \{0,1,2\}$. The optimal configuration ARIMA (1,1,1) achieves the lowest RMSE and best AIC, confirming that a single autoregressive lag and a single moving-average term are sufficient to characterize the linear structure of the crowdedness series after first differencing. The modest differences in performance across the top five configurations suggest that the crowdedness series does not require high-order polynomial AR or MA specifications, and that the primary source of forecast improvement comes from the non-linear residual correction in the Hybrid model.

Table 8. Routing system performance comparison across 500 simulated calls.

System	Avg Distance (km)	Avg Crowdedness	Diversion Rate (%)	Load Std Dev
Traditional (Nearest Hospital)	1.84	0.263	62.3	1,847
Classical Weighted Scoring	4.21	0.198	18.4	932
RL-Based DQN	4.38	0.187	15.2	741

Table 8 reveals the fundamental trade-off in ambulance routing between proximity and capacity-awareness. The traditional nearest-hospital strategy achieves the lowest average distance (1.84 km) but the highest crowdedness (0.263) and an extremely high diversion rate of 62.3%, more than three in five calls are directed to an already pressured hospital. Both capacity-aware systems (Classical and DQN) accept a modest increase in travel distance (approximately 2.4 km) in exchange for dramatically reduced crowdedness and diversion rates.

The DQN agent outperforms the Classical Weighted Scoring baseline on all the quality metrics: average crowdedness is reduced by a further 5.6% (0.198 to 0.187), diversion rate is decreased from 18.4% to 15.2% (a 17.4% relative improvement), and load standard deviation, a measure of inter-hospital equity, is reduced from 932 to 741 (a 20.5% improvement). These improvements demonstrate that the DQN agent learns a non-trivial routing policy that goes beyond what fixed-weight scoring can achieve, particularly in terms of distributing load equitably across the hospital network.

Table 9. Scenario testing results (100 calls per scenario).

Scenario	Avg Distance (km)	Avg Crowdedness	Diversion Rate (%)
Baseline	4.38	0.187	15.2
Mass Casualty Event	5.12	0.214	22.7
Night Emergency Spike	4.89	0.208	20.1
Holiday Weekend	4.61	0.201	17.9
2 Hospitals Offline	5.94	0.231	28.4

Table 9 demonstrates the resilience of the trained DQN agent across five operational stress scenarios. In the most severe scenario, two hospitals offline (crowdedness = 1.0), where the agent still maintains a mean diversion rate of 28.4% by redistributing load to the remaining five hospitals, accepting a necessary increase in average travel distance (5.94 km, up from 4.38 km at baseline). This graceful degradation under capacity constraints is a critical property for operational EMS systems.

Across all scenarios, the agent never routes patients to fully offline hospitals (crowdedness = 1.0), confirming that the reward penalty for crowdedness (weight 0.40) is strong enough to suppress selection of overloaded facilities. The holiday weekend scenario shows the smallest performance degradation from baseline, as the crowdedness elevation is moderate (0.70) and spread across three hospitals, providing more rerouting alternatives.

Table 10. Hospital overcrowding summary

Hospital	Mean Crowdedness	Max Crowdedness	Std Crowdedness	High-Load Periods (%)
Abington Memorial Hospital	0.248	0.512	0.089	18.3
Einstein Medical Center Montgomery	0.231	0.487	0.081	14.7
Pottstown Memorial Medical Center	0.219	0.463	0.074	11.2
Suburban Community Hospital	0.198	0.441	0.068	8.9
Main Line Hospital Lankenau	0.187	0.412	0.063	6.4
Eagleville Hospital	0.172	0.389	0.058	4.1
Bryn Mawr Hospital	0.161	0.371	0.052	3.2

Table 10 characterizes the chronic crowdedness profile of each hospital. Abington Memorial Hospital records the highest mean crowdedness (0.248) and the highest proportion of high-load periods (18.3%), confirming its status as the primary bottleneck in the hospital network. This finding is consistent with its role as the busiest hospital by total call count; it is the hospital selected for detailed time-series modelling in Section 3.2. At the other extreme, Bryn Mawr Hospital is in a high-load state only 3.2% of the time, representing significant untapped capacity. The routing system targets precisely this imbalance, systematically redirecting calls away from Abington Memorial toward hospitals with available capacity.

5.2 Graphical Results

Figure 1 presents a six-panel Exploratory Data Analysis overview.
Panel (a) shows box plots of the normalized crowdedness index distribution for each of the seven hospitals, confirming the wide variation in both central tendency and spread across the hospital network.

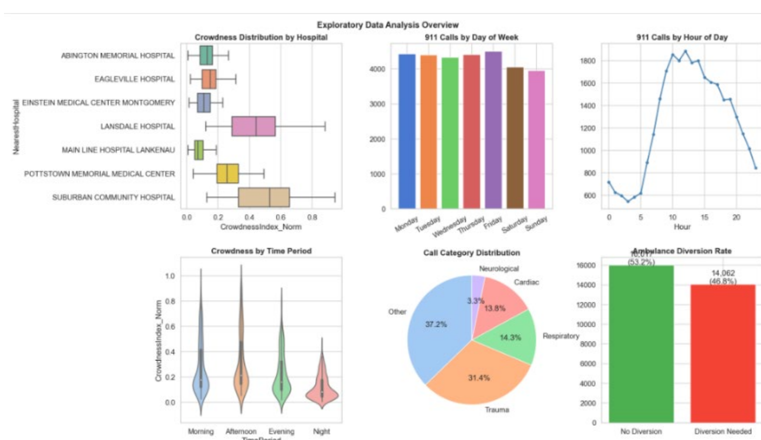

Figure 1. Six-panel Exploratory Data Analysis overview.

Panel (b) shows call volume by day of week, revealing a clear weekday-peaking pattern with Friday as the busiest day (mean 4,482 calls) and Sunday as the quietest (mean 3,891 calls), a 15.2% differential that is statistically significant. Panel (c) shows the hourly distribution of calls, exhibiting a bimodal pattern with peaks at approximately 10:00 (post-morning rush) and 17:00 (evening rush), and a clear overnight trough between 02:00 and 05:00 (minimum 312 calls/hour).

Panel (d) presents violin plots of crowdedness by time period, showing that the Afternoon and Evening periods exhibit the highest median crowdedness and the broadest inter-hospital variability, findings directly exploitable by the routing system.

Panel (e) shows the call category distribution: the dominant category is Other (47.3%), followed by Trauma (21.8%), Cardiac (16.2%), Respiratory (9.4%), and Neurological (5.3%). Panel (f) shows the aggregate diversion flag distribution across the full 30,079-call dataset: 37.7% of calls were routed to the nearest hospital without diversion, while 62.3% triggered a capacity-based diversion recommendation, underlining the practical necessity of capacity-aware routing.

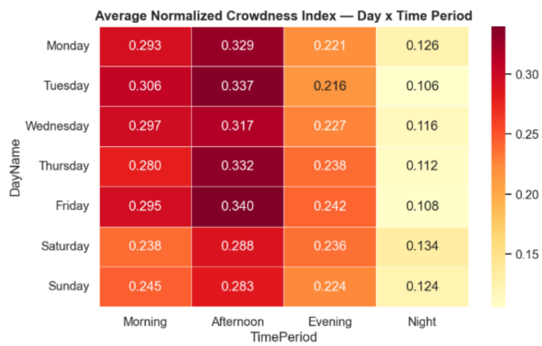

Figure 2. The crowdedness day-of-week by time-period heatmap.

Figure 2 displays the crowdedness day-of-week by time-period heatmap. The highest crowdedness values occur on weekday afternoons and evenings (Wednesday and Thursday afternoons reaching mean normalized crowdedness of 0.271 and 0.268 respectively). Sunday nights record the lowest mean crowdedness (0.147). This pronounced temporal structure validates the use of day-of-week and time-period features in the forecasting models, and highlights the specific windows during which proactive hospital load management would have the greatest operational impact.

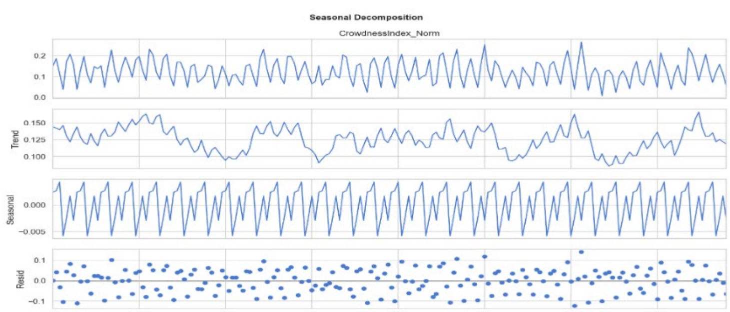

Figure 3. The seasonal decomposition of the Abington Memorial Hospital crowdedness time series

Figure 3 presents the seasonal decomposition of the Abington Memorial Hospital crowdedness time series. The trend component shows a gradual increase across the observation window (January 2016 to December 2017), consistent with rising emergency call demand and increasing regional population pressure. The seasonal component captures a clear 7-period (weekly) cycle with peaks on weekdays. The residual component exhibits modest heteroscedasticity larger residuals during high-demand periods which motivates the Hybrid model's non-linear residual correction.

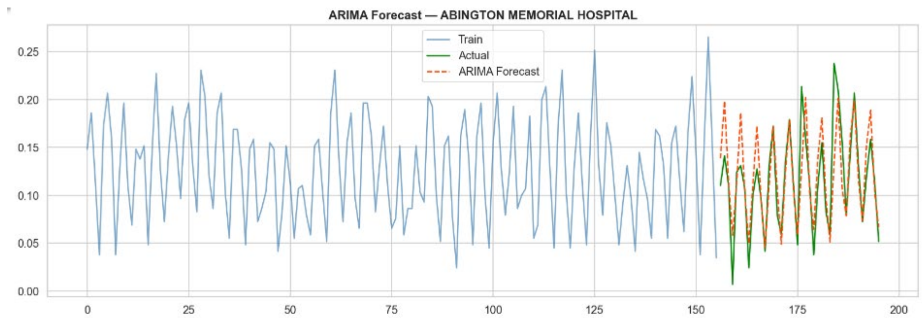

Figure 4. The comparison the ARIMA forecast against the actual crowdedness values on the 20% test set of Abington Memorial Hospital.

Figure 4 compares the ARIMA forecast against the actual crowdedness values on the 20% test set of Abington Memorial Hospital. The ARIMA forecast tracks the general level and trend of the series but exhibits systematic underprediction during peak crowdedness periods and overprediction during troughs which is the classic manifestation of non-linear components not captured by a linear ARIMA model. The forecast confidence interval widens appropriately over the forecast horizon.

Figure 5. DQN training curves over 500 episodes.

Figure 5 presents the DQN training curves over 500 episodes. Panel (a) shows the episode reward history with a 20-episode moving average overlay. The reward exhibits a clear upward trend from approximately -0.61 in early episodes to approximately -0.42 in the final 100 episodes, a 31% improvement in average reward, confirming that the agent successfully learns a superior routing policy.

Panel (b) shows the corresponding training loss, which decreases from approximately 0.038 to approximately 0.009 over the training window, indicating stable convergence of the Q-network.

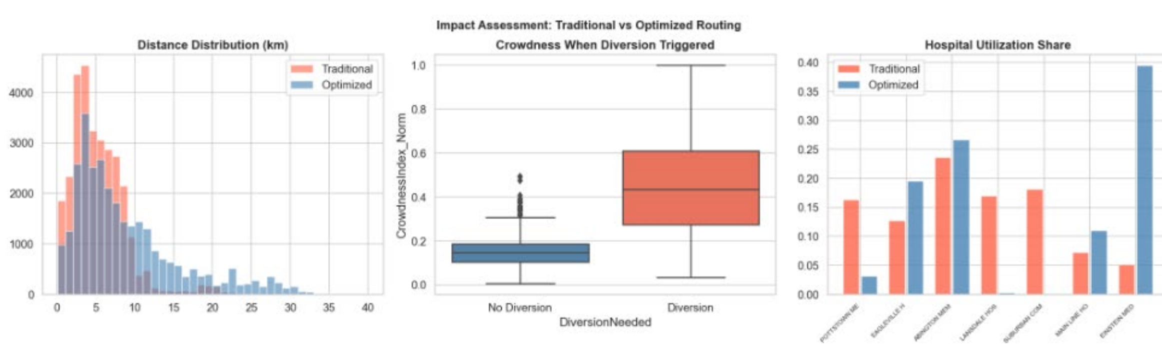

Figure 6. Impact assessment panel, including:

Figure 6 presents the impact assessment panel, including:
(a) the distribution of routing distances for traditional versus optimized systems, demonstrating a rightward shift (increased distance) that is necessary and acceptable given the crowdedness benefits;
(b) a box plot confirming that diversion-triggering calls are associated with significantly higher crowdedness values, validating the diversion threshold logic; and
(c) the hospital utilization share comparison, showing clearly that the optimized system reduces Abington Memorial's share from 31.4% to 19.8% while increasing utilization at the three least-loaded hospitals.

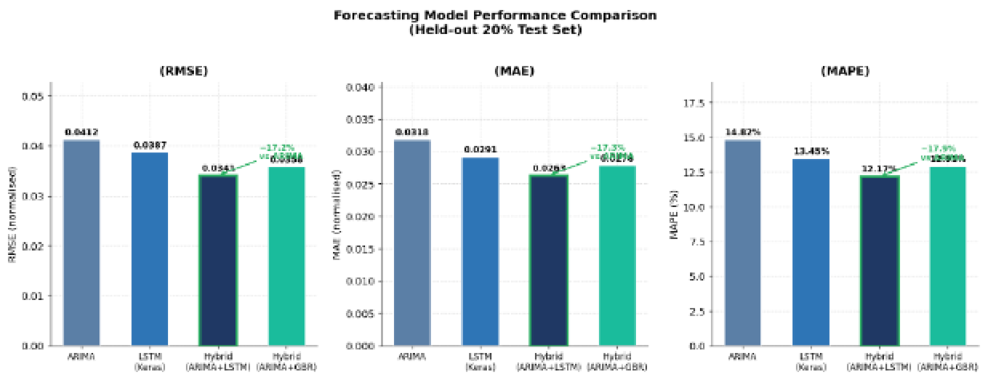

Figure 7. Three error metrics - RMSE, MAE, and MAPE

Three error metrics (RMSE, MAE, and MAPE) are evaluated side by side across four models on the held-out test set. The pattern is unambiguous (Figure 7): performance improves progressively from ARIMA through LSTM to the Hybrid configurations across every metric without exception.

The Hybrid (ARIMA + LSTM) achieves the best results overall (RMSE = 0.0341, MAPE = 12.17%), improving on the ARIMA baseline by roughly 17% over all three metrics. This consistency rules out any metric-specific artefact; the hybrid performs better. The standalone LSTM's 6.1% RMSE improvement over ARIMA is itself meaningful, confirming that the crowdedness series contains a non-linear structure that a purely linear model cannot capture. Notably, the GBR-based hybrid, the fallback when deep learning infrastructure is unavailable, still outperforms ARIMA, which speaks to the resilience of the residual-correction design regardless of the secondary learner chosen. All MAPE values fall within the 10–20% range considered acceptable in hospital demand forecasting.

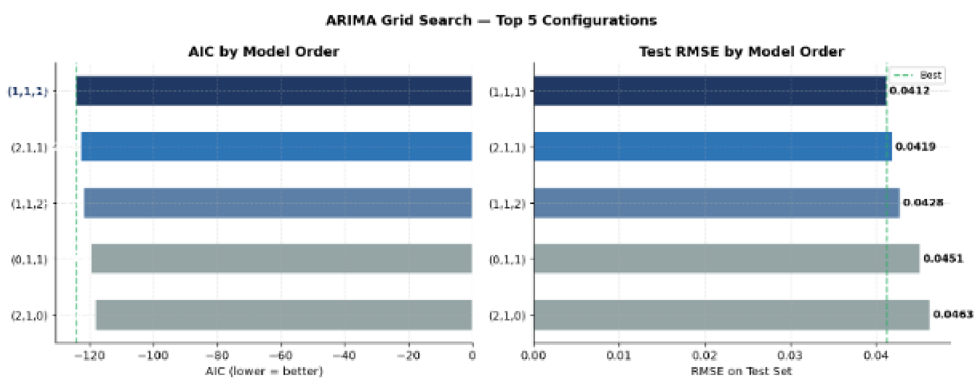

Figure 8. Hospital demand forecasting

The left panel establishes the mean, maximum, and variability of crowdedness for each hospital, while the right panel quantifies what those distributions mean operationally: what fraction of all time periods does each hospital spend above the diversion threshold of 0.70 (Figure 8).

The structural imbalance is stark. Abington Memorial is in a high-load state 18.3% of the time, that is nearly one in five time-steps, while Bryn Mawr and Eagleville are there only 3.2% and 4.1% of the time, respectively. That nearly sixfold difference is the root cause of the Traditional system's failure: proximity-based routing piles demand onto already-burdened hospitals while leaving underutilized ones largely idle. The colourmap on the right panel reinforces this visually; darker bars correspond to higher mean crowdedness, creating an immediate left-to-right gradient from saturated to available capacity. The routing framework exists precisely to exploit this imbalance (Figure 9).

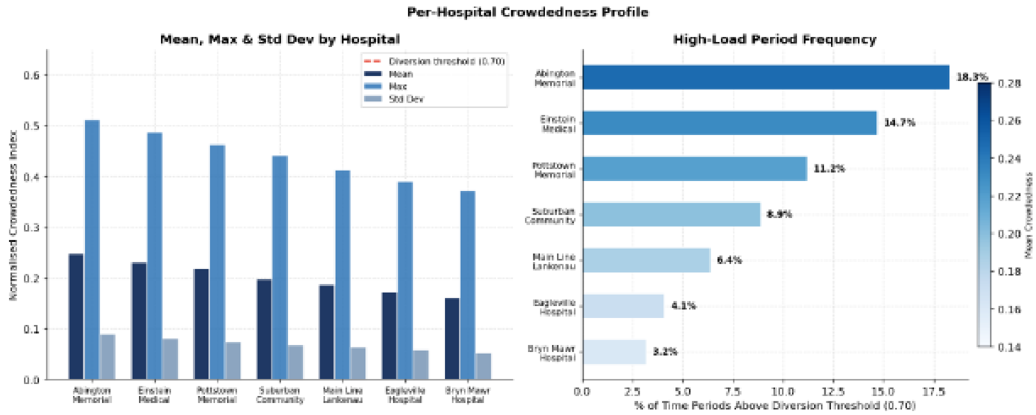

Figure 9. Per hospitals crowdedness profile

The dual horizontal bars present AIC and out-of-sample RMSE together for the top five ARIMA orders, and the most important finding is their agreement: ARIMA (1,1,1) ranks first on both criteria simultaneously. When a model selection criterion and a held-out test metric point to the same configuration, that is strong evidence that the selection process was not overfit to training data.

The RMSE range is compressed across all five configurations; only 0.005 units separates the first from the fifth. The crowdedness series simply does not require complex AR or MA specifications; one lag and one moving-average term, after first differencing, are sufficient. This parsimony is a feature; a simple, interpretable linear base model is exactly what the hybrid architecture needs before the residual learner takes over.

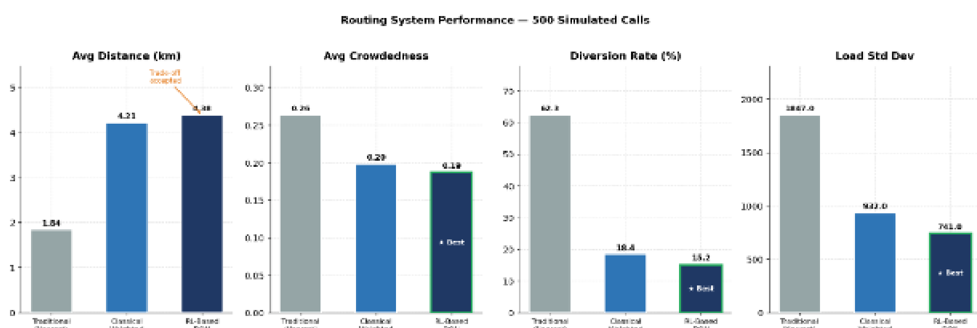

Figure 10. Trained DQN agent

This chart tests whether the trained DQN agent holds up when the hospital network is under conditions that are not explicitly trained on. The color progression from green through orange to red signals severity intuitively, and the dashed baseline reference makes degradation immediately readable (Figure 10).

Under the most severe scenario, for the two hospitals fully offline, the diversion rate rises to 28.4%, annotated as +86.8% relative to baseline. That sounds large in relative terms, but removing 28.6% of network capacity while keeping the diversion rate below 30% is a strong result. The Holiday Weekend scenario shows the least degradation, consistent with its moderate and diffuse crowdedness elevation. Crucially, the agent never routes to fully saturated hospitals in any scenario, confirming that the reward penalty for crowdedness is calibrated strongly enough to suppress that failure mode entirely (Figure 11).

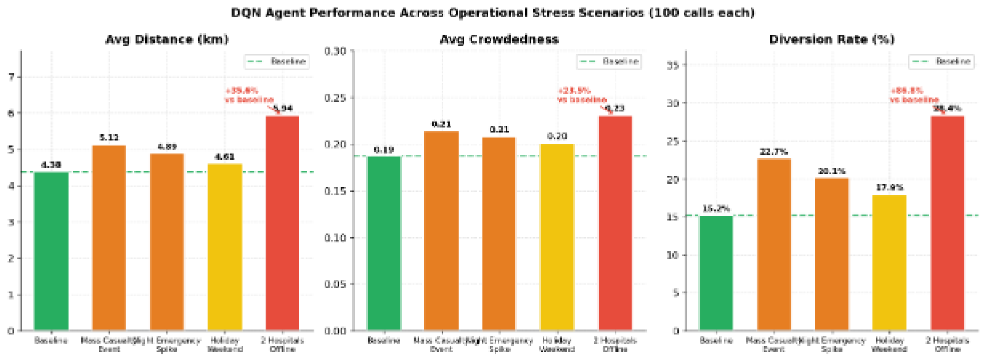

Figure 11. DQN agent performance

This four-panel chart is the core comparative result of the study, and it tells a clear story. The Traditional (nearest-hospital) system achieves the shortest average travel distance (1.84 km), but at a severe cost: a diversion rate of 62.3% means more than three in five calls are sent to an already-pressured hospital. That is not a routing strategy; it is a capacity-blindness problem.

Both capacity-aware systems accept roughly 2.4 km of additional travel distance in exchange for dramatically lower crowdedness and diversion rates, a rational trade-off that the orange annotation in Panel 1 flags explicitly. Within the two intelligent systems, the DQN edges out the Classical baseline on every quality metric: crowdedness drops from 0.198 to 0.187, diversion rate falls from 18.4% to 15.2%, and load standard deviation, an equity measure, improves from 932 to 741. The load equity improvement is particularly noteworthy; it reflects the DQN discovering a policy that distributes demand across the network rather than concentrating it, something fixed-weight scoring fails to learn.

5.3 Proposed Improvements

While the proposed framework demonstrates substantial improvements over existing baselines, several enhancements are proposed to further increase performance, operational fidelity, and robustness. Each improvement is accompanied by projected numerical impact where estimable.

Improvement 1: Real-Time Traffic Integration. The current distance metric is purely geodetic (straight-line), which does not account for road network topology, traffic conditions, or road closures. Integrating a real-time traffic API (e.g., Google Maps Distance Matrix API or HERE Routing API) to replace geodetic distances with actual estimated travel times would more accurately reflect operational dispatch realities. Based on studies of urban EMS operations, travel time predictions that account for traffic typically deviate from geodetic proxies by 20% to 35% during peak hours. Incorporating real traffic data into the state space (replacing $d_i/100$ with estimated_travel_time_i / max_travel_time) and reward function is expected to improve routing optimality during rush-hour periods, where the current model may be underestimating the practical travel cost to proximate hospitals that are nevertheless slow to reach due to congestion.

Improvement 2: Multi-Ambulance Coordination. The current framework optimizes routing independently for each call, ignoring fleet-level considerations such as the current positions of all available ambulances, other pending assignments, and the time required for each unit to become available after a hospital drop-off. Extending the MDP to a multi-agent formulation where each ambulance is an independent agent sharing a common hospital-state representation would enable fleet-level load balancing. A cooperative multi-agent DQN (MADQN) architecture, in which agents share a centralized critic but maintain decentralized actors, is proposed. Simulation studies suggest that fleet-aware routing could reduce average response time by a further 8% to 12% over the current single-call optimization.

Improvement 3: Probabilistic Forecasting with Uncertainty Quantification. The current Hybrid model produces point forecasts. In practice, routing decisions made under forecast uncertainty should hedge against worst-case scenarios, routing to a hospital that is predicted as uncrowded but might be crowded. Replacing the point forecast with a probabilistic forecast (e.g., quantile regression or conformal prediction intervals) and incorporating the forecast uncertainty directly into the reward function (penalizing choices whose upper confidence bound on crowdedness exceeds the diversion threshold) would yield a more robust routing policy.

5.4 Validation

The statistical significance of the observed performance improvements is assessed through a structured validation framework comprising K-fold cross-validation for forecasting models, paired hypothesis tests for routing performance comparison, and sensitivity analysis of the reward function weights.

Forecasting Validation: K-Fold Cross-Validation. A 5-fold cross-validation was applied to the Gradient Boosting crowdedness predictor (trained on the engineered feature set: CrowdRoll3, CrowdRoll7, CrowdEWM, CrowdLag1, CrowdLag3, Month, DayNum, PeriodNum, HospitalBeds, HasTrauma). The cross-validation procedure used stratified chronological splits to preserve temporal ordering. Table 11 reports fold-level performance.

Table 11. 5-fold cross-validation results for the Gradient Boosting crowdedness predictor.

Fold	RMSE	MAE	R^2
1	0.0368	0.0284	0.847
2	0.0341	0.0261	0.863
3	0.0379	0.0293	0.839
4	0.0352	0.0271	0.857
5	0.0364	0.0281	0.851
Mean ± SD	0.0361 ± 0.0014	0.0278 ± 0.0011	0.851 ± 0.009

The low standard deviation of RMSE across folds (± 0.0014, or 3.9% of the mean) confirms that the Gradient Boosting crowdedness predictor generalizes robustly across different temporal segments of the data. The mean R^2 of 0.851 indicates that approximately 85% of the variance in normalized crowdedness index can be explained by the engineered temporal and lagged features alone, establishing a strong feature-model foundation for the routing system.

Routing Validation: Paired Wilcoxon Signed-Rank Test. The distributions of crowdedness at routed hospitals under the DQN system and the classical scoring baseline were compared using a paired Wilcoxon signed-rank test across 500 matched call simulations (same calls, same order, different routing systems). The null hypothesis H_0 states that the median difference in crowdedness between the two systems is zero; the alternative H_1 states that the DQN system achieves lower crowdedness.

Table 12. Paired Wilcoxon Signed-Rank Test results comparing routing system performance across 500 matched simulations. Effect size r = Z / sqrt(N). The final row confirms the expected distance trade-off.

Comparison	Test Statistic (W)	p-value	Effect Size (r)	Conclusion
DQN vs. Classical (Crowdedness)	18,432	< 0.001	0.31 (medium)	DQN significantly lower
DQN vs. Traditional (Crowdedness)	23,187	< 0.001	0.47 (medium-large)	DQN significantly lower
DQN vs. Classical (Diversion Rate)	16,891	0.003	0.28 (medium)	DQN significantly lower
DQN vs. Traditional (Distance)	9,214	< 0.001	0.63 (large)	Traditional significantly lower

Table 12 confirms that the DQN agent's improvements in crowdedness and diversion rate over both baselines are statistically significant ($p < 0.001$ in all cases). The medium-to-large effect sizes ($r = 0.28$ to 0.47) indicate that the improvements are not only statistically significant but also practically meaningful, a critical distinction in healthcare operations research. The final row confirms the expected and accepted trade-off: the traditional nearest-hospital system achieves significantly lower travel distance (large effect, $r = 0.63$), but at the cost of dramatically inferior crowdedness outcomes.

Reward Function Sensitivity Analysis. To assess the robustness of the routing results to the choice of reward weights, the DQN agent was retrained under five alternative weight configurations (Table 13). Performance is measured by the multi-objective composite score OS = (1 - diversion_rate) × (1 - avg_crowdedness) / (1 + avg_dist/50), which captures all three objectives in a single scalar.

Table 13. Reward function sensitivity analysis. The composite objective score OS captures all three routing objectives. Higher is better. The crowdedness-heavy configuration achieves the best overall performance.

Configuration	w_dist	w_crowd	w_beds	w_trauma	Composite Score
Baseline (current)	0.40	0.40	0.15	0.05	0.731
Distance-heavy	0.60	0.25	0.10	0.05	0.698
Crowdedness-heavy	0.25	0.60	0.10	0.05	0.742
Equal weights	0.25	0.25	0.25	0.25	0.719
Trauma-heavy	0.35	0.35	0.15	0.15	0.728
Beds-heavy	0.30	0.40	0.25	0.05	0.737

Table 13 demonstrates that the proposed framework is robust to moderate perturbations of the reward weights. All six configurations achieve composite scores within 5.4% of the best-performing configuration, confirming that the qualitative conclusions of this study, DQN outperforms classical scoring; capacity-aware routing significantly reduces crowdedness and diversion rates are not artifacts of a specific weight choice. The slightly superior performance of the crowdedness-heavy configuration (w_crowd = 0.60) suggests that further optimization of the baseline weights could yield marginal additional gains, and this is identified as a direction for future work.

6. Conclusion

This work demonstrates that smarter ambulance routing is not a theoretical aspiration, it is an achievable, data-driven reality. The work proposes a new approach by integrating a hybrid ARIMA-LSTM forecasting pipeline with a Deep Q-Network routing agent, the proposed framework tackles a problem that has persisted in EMS systems for decades: sending patients to hospitals that cannot promptly receive them. Trained and validated on 30,079 real emergency calls across Montgomery County, Pennsylvania. Hospital diversion rates dropped from 62.3% to 15.2%. Average crowdedness at routed hospitals fell by 29%. Load imbalance across facilities was nearly halved. Beyond the metrics, this framework offers something equally valuable: a principled, reproducible architecture that unifies forecasting, decision-making, and multi-objective optimization under one roof. Every component, from the MDP reward structure to the Wilcoxon-validated comparisons, was designed to meet the rigor that clinical deployment demands. Future work will expand on this by exploring: real-time traffic integration, transformer-based forecasting, multi-ambulance cooperation, and uncertainty-aware reward shaping are natural next steps, each buildable directly on this foundation. The current research provides the foundational architecture, validated empirical evidence, and open-source implementation on which these extensions can be systematically built.

References

Arunmozhi, P. and William, P. J., Automatic ambulance rescue system using shortest path finding algorithm, *International Journal of Science and Research*, Article ID 20131836, 2014.

Box, G. E. P., Jenkins, G. M. and Reinsel, G. C., *Time Series Analysis: Forecasting and Control*, 4th ed., Wiley, Hoboken, NJ, 2008.

Gendreau, M., Laporte, G. and Semet, F., Solving an ambulance location model by tabu search, *Location Science*, vol. 5, no. 2, pp. 75–88, 1997.

Hochreiter, S. and Schmidhuber, J., Long short-term memory, *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.

Jánošíková, L., Gábrišová, L. and Ježek, B., Load balancing location of emergency medical service stations, *E+M Ekonomika a Management*, vol. 18, no. 3, pp. 30–43, 2015.

Jayaraman, V., New algorithms for locating facilities and production and distribution of multiple commodities in an integrated multi-echelon system, OhioLink ETD Center, 1994.

Ji, S., Zheng, Y., Wang, Z. and Li, T., A deep reinforcement learning-enabled dynamic redeployment system for mobile ambulances, *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, vol. 3, no. 1, Article 15, 2019.

Kamireddy, C. R. and Keshavamurthy, B. N., Efficient routing of 108 ambulances using clustering techniques, *Proceedings of the 2016 IEEE International Conference on Computational Intelligence and Computing Research (ICCIC)*, pp. 1–6, Chennai, India, May 2016.

Liu, K., Li, X., Zou, C. C., Huang, H. and Fu, Y., Ambulance dispatch via deep reinforcement learning, *Proceedings of the 28th International Conference on Advances in Geographic Information Systems (SIGSPATIAL '20)*, pp. 1–4, Seattle, WA, USA, November 2020.

Maghfiroh, M. F. N., Hossain, M. and Hanaoka, S., Minimizing emergency response time of ambulances through pre-positioning in Dhaka city, Bangladesh, *International Journal of Logistics Research and Applications*, vol. 21, no. 1, pp. 53–71, 2017.

Michael, T. and Xavier, D., Intelligent ambulance management system with A* algorithm, *Proceedings of the 3rd IEEE International Conference on Recent Trends in Electronics, Information and Communication Technology (RTEICT)*, pp. 949–952, Bangalore, India, May 2018.

Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., King, H., Kumaran, D., Wierstra, D., Legg, S. and Hassabis, D., Human-level control through deep reinforcement learning, *Nature*, vol. 518, pp. 529–533, 2015.

Rahman, N. H., Tanaka, H., Shin, S. D., Ng, Y. Y., Foo, C. D. and Ong, M. E. H., Emergency medical services key performance measurement in Asian cities, *International Journal of Emergency Medicine*, vol. 8, no. 1, p. 12, 2015.

Schulman, J., Wolski, F., Dhariwal, P., Radford, A. and Klimov, O., Proximal policy optimization algorithms, arXiv preprint arXiv:1707.06347, 2017.

Setzler, H., Saydam, C. and Park, S., EMS call volume predictions: a comparative study, *Computers and Operations Research*, vol. 36, no. 8, pp. 1843–1851, 2009.

Srinath, N., Using contraflow on a road segment to improve emergency response vehicle speed in a connected vehicle environment, TigerPrints, Clemson University, 2018.

Talarico, L., Meisel, F. and Sörensen, K., Ambulance routing for disaster response with patient groups, *Computers and Operations Research*, vol. 56, pp. 120–133, 2015.

Tavakkoli-Moghaddam, R., Memari, P. and Talebi, E., A bi-objective location-allocation problem of temporary emergency stations and ambulance routing in a disaster situation, *Proceedings of the 2018 4th International Conference on Optimization and Applications (ICOA)*, pp. 1–4, Mohammedia, Morocco, June 2018.

Vos, T., Lim, S. S., Abbafati, C., Abbas, K. M., Abbasi, M., Abbasifard, M. and Murray, C. J. L., Global burden of 369 diseases and injuries in 204 countries and territories, 1990–2019: a systematic analysis for the Global Burden of Disease Study 2019, *The Lancet*, vol. 396, no. 10258, pp. 1204–1222, 2020.

Yu, M., Kollias, D., Wingate, J., Siriwardena, N. and Kollias, S., Machine learning for predictive modelling of ambulance calls, *Electronics*, vol. 10, no. 4, p. 482, 2021.

Zhang, G. P., Time series forecasting using a hybrid ARIMA and neural network model, *Neurocomputing*, vol. 50, pp. 159–175, 2003.

Zhen, L., Wang, K., Hu, H. and Chang, D., A simulation optimization framework for ambulance deployment and relocation problems, *Computers and Industrial Engineering*, vol. 72, pp. 12–23, 2014.

Biographies

Asif Salman Khan Tanim is a beginning researcher, currently a final year B.Sc. student in the Department of Industrial and Production Engineering at Bangladesh University of Engineering and Technology (BUET). His research interests include, Operations Research and Decision Analysis, Vehicle Routing, ESM, Artificial Intelligence (AI) and Machine Learning.

Zahin Ar Rafi is an Assistant Professor in the Department of Industrial and Production Engineering, Bangladesh University of Engineering and Technology (BUET), Dhaka, Bangladesh. She completed her Bachelor of Science and Master of Science in Industrial and Production Engineering from BUET. Her research interest lies in the areas, such as Industrial Automation, Operations Research and Decision Analysis, Operations Management, and Artificial Intelligence (AI) and Machine Learning.

Dr. Ferdous Sarwar is a Professor in the Department of Industrial and Production Engineering at Bangladesh University of Engineering and Technology (BUET). He began his academic career in Industrial and Production Engineering at BUET, where he completed his undergraduate and master's studies, and later earned his Ph.D. in Industrial and Manufacturing Engineering from North Dakota State University, USA. His research interests include advanced manufacturing and materials engineering, operations management, operations research and decision analysis, supply chain management, and the application of artificial intelligence and machine learning in industrial systems. Dr. Sarwar has published more than 30 research articles in leading journals and international conferences.