Simen Owen Academic Proceedings Series

3rd International Conference on Electronics, Engineering, Computer Science and Applied Development (EESD 2026)

Article

Time-Series Demand Forecasting and Resource Allocation Decision Optimization Model for Smart Healthcare Systems

Chenyuan Xu [1,*]

[1] College of Science, Mathematics and Technology, Wenzhou-Kean University, Wenzhou, China
* Correspondence: Chenyuan Xu, College of Science, Mathematics and Technology, Wenzhou-Kean University, Wenzhou, China

Abstract: To address the critical spatiotemporal imbalance of medical resource supply and demand, as well as the pervasive issue of system congestion in modern smart healthcare environments, this paper proposes a comprehensive closed-loop management model that seamlessly integrates demand time-series forecasting with resource allocation decision optimization. First, utilizing extensive patient admission logs and detailed bed utilization records extracted from the widely recognized MIMIC-IV clinical database, we construct robust, multi-dimensional time-series feature engineering. A sophisticated combined XGBoost machine learning model is subsequently employed to accurately predict medical service demand for future operational periods. This predictive framework effectively captures the complex, non-linear fluctuation patterns characteristic of patient arrival rates in dynamic clinical settings. Second, based on these highly accurate prediction results, a rigorous bi-objective resource scheduling optimization model is established. This model is specifically designed to simultaneously minimize patient waiting costs and maximize overall resource utilization rates across various hospital departments. Furthermore, a dynamic recommendation mechanism is introduced to automatically generate optimal scheduling and bed allocation schemes tailored to real-time conditions. Finally, comprehensive simulation experiments conducted using the MIMIC-IV demo data conclusively demonstrate that, when compared with traditional fixed allocation modes, the proposed integrated model significantly reduces patient queue waiting times. Moreover, it substantially improves the overall operational efficiency and adaptability of medical resources, thereby providing vital, data-driven decision support for the refined, sustainable management of next-generation smart hospitals.

Keywords: smart healthcare; time-series forecasting; resource allocation; decision optimization; machine learning

Received: 07 April 2026
Revised: 21 May 2026
Accepted: 01 June 2026
Published: 05 June 2026

1. Introduction

The escalating demand for healthcare services, driven by aging populations and the increasing prevalence of chronic diseases, has placed unprecedented pressure on medical infrastructure worldwide [1]. A critical challenge in modern hospital management is the stochastic nature of patient inflow, which frequently mismatches with the fixed supply of medical resources such as emergency department staff, inpatient beds, and diagnostic equipment. This supply-demand imbalance manifests as severe system congestion, prolonged patient waiting times, and suboptimal resource utilization, ultimately compromising the quality of care and patient safety. In the context of "Smart Healthcare," the widespread adoption of Hospital Information Systems (HIS) has generated vast repositories of operational data, including patient arrival timestamps, service duration logs, and equipment usage records. Consequently, shifting from reactive, experience-based management to proactive, data-driven decision-making has become a focal point for improving healthcare service efficiency.

Academic and industrial efforts to address these challenges have largely evolved along two distinct but complementary trajectories: demand forecasting and resource allocation optimization. In the domain of demand forecasting, methodologies have transitioned from traditional linear statistical models to advanced machine learning techniques capable of handling high-dimensional data. Contemporary research indicates that patient arrival rates exhibit complex time-series characteristics, including seasonality, trend, and non-linear volatility [2]. Studies have demonstrated that incorporating multi-dimensional features, such as calendar effects, meteorological data, and historical lag variables, significantly enhances prediction accuracy. Furthermore, recent advancements suggest that integrating external data sources, such as internet search indices, with nonlinear models like neural networks or tree-based ensemble algorithms can better capture the dynamic fluctuations of medical demand compared to static baselines.

However, accurate forecasting alone is insufficient to resolve systemic congestion; it must be effectively translated into actionable operational strategies. In the realm of resource allocation, optimization models often employ discrete event simulation and heuristic algorithms to solve complex scheduling problems. The primary objective in this field is to balance conflicting goals, specifically minimizing patient waiting costs while maximizing resource utilization rates [3]. Despite significant progress, a disconnect remains in many existing frameworks where prediction models operate independently of decision-support systems, or optimization algorithms rely on static assumptions rather than dynamic, real-time demand signals. This fragmentation limits the ability of healthcare systems to respond agility to sudden surges in patient volume.

To bridge this gap, this dissertation proposes a comprehensive model for time-series demand forecasting and resource allocation decision optimization within smart healthcare systems. By leveraging granular clinical operation data, this research first constructs a feature engineering framework to capture the temporal patterns of patient flow and resource usage. A high-precision hybrid prediction model is then developed to forecast short-term service demand [4]. Crucially, these predictions serve as dynamic inputs for a bi-objective resource configuration optimization model. A recommendation mechanism is subsequently introduced to generate optimal staffing and equipment scheduling schemes. This closed-loop approach aims to mathematically minimize system congestion and enhance service efficiency, providing a robust theoretical and practical foundation for the intelligent management of modern medical facilities.

2. Theoretical Framework

2.1. Limitations of Traditional Time-Series Forecasting in Healthcare

The prediction of patient demand within healthcare systems has traditionally relied on linear statistical models. While these autoregressive approaches offer interpretability and computational simplicity, their rigid structural assumptions often fail to capture the high-frequency volatility and non-linear complexities inherent in medical operational data [4]. Healthcare demand is rarely stationary; it is subject to stochastic surges driven by seasonal epidemics, public health events, and complex calendar effects. Consequently, models that presume linearity or stationarity frequently exhibit substantial lag errors when confronting sudden shifts in patient arrival rates or equipment usage frequency. Although recent advancements have introduced multivariate extensions to incorporate external variables, the fundamental inability of these static statistical frameworks to model long-term dependencies and intricate feature interactions limits their utility in dynamic, high-load smart healthcare environments.

2.2. The Shift to Machine Learning and Feature Engineering Challenges

In response to the inadequacies of statistical baselines, the academic focus has shifted toward machine learning (ML) and deep learning paradigms. Algorithms such as gradient boosting decision trees and recurrent neural networks have demonstrated superior capability in discerning latent patterns within high-dimensional datasets. However, a critical examination reveals that the performance of these "black-box" models is heavily contingent upon the quality of feature engineering. Existing literature often oversimplifies the temporal input, neglecting the significance of granular "lag features," "rolling window statistics," and multi-source external determinants (e.g., meteorological data or internet search indices). Furthermore, while ML models excel in minimizing error metrics (such as RMSE), a pure pursuit of accuracy often ignores the practical interpretability required for clinical decision-making [4]. A model that predicts a surge in demand but fails to elucidate the contributing factors offers limited actionable value to hospital administrators.

2.3. Critical Analysis of Resource Allocation and Optimization Strategies

Parallel to forecasting, resource allocation strategies have evolved from heuristic, experience-based scheduling to sophisticated mathematical optimization. The integration of Discrete Event Simulation (DES) with meta-heuristic algorithms, such as Genetic Algorithms or Particle Swarm Optimization, represents the current state-of-the-art. These methods attempt to solve the multi-objective problem of minimizing patient waiting times while maximizing resource utilization.

However, a significant methodological flaw persists in many existing optimization frameworks: the reliance on deterministic or retrospective data. Optimization models often treat patient demand as a known, static parameter or a simple average of historical data, rather than a dynamic, probabilistic variable derived from real-time predictions. This static assumption renders the resulting resource configuration schemes brittle; an "optimal" schedule derived from average historical data often collapses under the pressure of real-world variance. Additionally, the computational complexity of simulation-based optimization frequently precludes real-time application, creating a latency barrier between data analysis and operational execution.

2.4. The Disconnect between Prediction and Decision Support

Perhaps the most critical gap identified in the current theoretical landscape is the systemic disconnect between predictive analytics and prescriptive decision support. The majority of research efforts treat demand forecasting and resource optimization as isolated silos. Prediction studies typically conclude with the validation of model accuracy, failing to demonstrate how these forecasts can translate into tangible efficiency gains. Conversely, operations research studies often assume perfect foresight of demand, neglecting the uncertainty inherent in predictive inputs [5].

True "smart healthcare" requires a cohesive, closed-loop mechanism where the output of the time-series prediction model serves as the direct, dynamic input for the optimization algorithm. The absence of an integrated "prediction-optimization-recommendation" framework limits the ability of current systems to proactively mitigate congestion. Therefore, establishing a methodology that bridges this divide, transforming high-precision forecasts into robust, implementable resource allocation recommendations, remains an imperative and under-explored frontier in healthcare management engineering [6].

3. Multi-Dimensional Time-Series Demand Prediction Model

3.1. Overview of the Prediction Framework

The foundational component of the proposed smart healthcare management system is the accurate forecasting of medical service demand. Unlike commercial supply chains, healthcare demand, specifically patient arrival rates and equipment usage, exhibits high stochasticity, strong seasonality, and complex non-linear dependencies. Traditional linear models often fail to capture abrupt fluctuations inherent in emergency department (ED) or intensive care unit (ICU) workflows. To address these challenges, this study constructs a data-driven prediction framework that integrates data ingestion, feature extraction, and model training [7]. The overall architecture of this framework is illustrated in Figure 1. As depicted, the pipeline transforms raw transactional logs from the Hospital Information System (HIS), such as the MIMIC-IV database, into structured time-series datasets. This end-to-end design ensures that raw operational data is efficiently converted into high-fidelity demand signals necessary for downstream resource optimization.

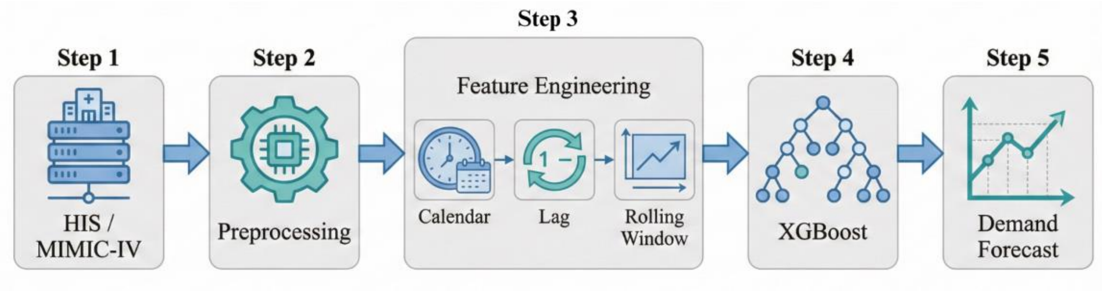

Figure 1. Overall Architecture of the Data-Driven Demand Prediction Framework.

3.2. Mathematical Formulation of the Problem

The demand prediction problem is formalized as a supervised learning task on a multivariate time series. The objective is to learn a mapping function that minimizes the expected loss over a future horizon. Specifically, the model seeks to predict the demand at a future time step based on a history of observed features [8]. This relationship can be expressed as:

$$\hat{y}_{t+h} = f(x_t, x_{t-1}, \dots, x_{t-k}; \theta) + \epsilon_t \quad (1)$$

In this formulation, $\hat{y}_{t+h}$ denotes the predicted demand at future time $t+h$, while x_{t-k} represents historical feature vectors up to lag k. The parameter θ embodies the trainable weights, and ϵ_t represents the irreducible error [9]. The learning algorithm optimizes θ such that the difference between the predicted value and the actual observed demand is minimized across the validation set, treating temporal dependencies through the structure of input features.

3.3. Data Preprocessing and Feature Engineering

Raw medical event logs are often sparse and noisy. To transform these discrete events into a continuous time series suitable for regression, a rigorous preprocessing pipeline aggregates timestamps into fixed intervals (e.g., hourly bins) and applies moving average smoothing to mitigate outliers. Building upon this, feature engineering is critical for capturing complex dynamics [10]. Three categories of features are constructed: Lag Features to capture autoregressive dependencies (e.g., demand at $t-1$); Calendar Features (hour, day, month) encoded via sine/cosine transformations to preserve cyclical continuity; and Rolling Window Statistics (mean, variance) calculated over sliding windows to provide context on recent volatility. This multi-dimensional feature space enables the algorithm to distinguish between routine fluctuations and significant structural shifts in demand.

3.4. Gradient Boosting Prediction Architecture

To handle non-linear interactions between features, this study employs the eXtreme Gradient Boosting (XGBoost) algorithm. Unlike linear ARIMA methodologies, XGBoost is an ensemble method that constructs a robust predictor by combining multiple weak decision trees additively [11]. The objective function combines a convex loss function, measuring prediction error, with a regularization term Ω that penalizes model complexity to prevent overfitting.

$$\mathcal{Obj}(\theta) = \sum_{i=1}^{n} l(y_i, \hat{y}_i) + \sum_{k=1}^{K} \Omega(f_k) \quad (2)$$

The regularization term is crucial in medical contexts where data is noisy, controlling tree complexity by penalizing the number of leaves and leaf weights. By optimizing this objective, the model balances accuracy with generalization, effectively capturing non-linear relationships such as sudden exponential rises in patient admissions during epidemic outbreaks [12].

3.5. Model Evaluation and Validation

To rigorously assess performance, the dataset is partitioned chronologically into training, validation, and testing sets to prevent data leakage [13]. The model is evaluated using Root Mean Square Error (RMSE) and Mean Absolute Percentage Error (MAPE).

$$\mathrm{RMSE} = \sqrt{\frac{1}{N}\sum_{t=1}^{N}(y_t - \hat{y}_t)^2} \tag{3}$$

RMSE is sensitive to large errors, making it suitable for detecting significant deviations in peak demand, while MAPE provides a scale-independent, interpretable measure of accuracy [14]. Minimizing these metrics ensures that the subsequent resource allocation module receives high-fidelity input signals, confirming the framework as a practical tool for driving the optimization engine.

4. Prediction-Driven Resource Allocation and Recommendation Mechanism

4.1. Transition from Prediction to Decision

High-precision demand forecasting, as established in Chapter 3, provides a necessary but insufficient condition for resolving healthcare congestion. The operational value of a prediction model is realized only when it drives effective decision-making [15]. In a typical hospital setting, resource managers face a fundamental trade-off: over-provisioning leads to wasted capacity and high operational costs, while under-provisioning results in dangerous overcrowding, long waiting times, and potential adverse medical outcomes. To address this, this chapter proposes a dynamic resource allocation framework. By treating the predicted demand $\hat{y}_t$ as a deterministic input, we construct a bi-objective optimization model designed to determine the optimal configuration of medical resources, such as staffing levels and active ICU beds, for future time windows.

4.2. Mathematical Formulation of the Optimization Model

The core of the decision support system is formulated as a discrete optimization problem. Let $t \in \{1, \ldots, T\}$ represent the planning horizon (e.g., the next 24 hours). Let x_t denote the quantity of active resources allocated at time t (e.g., the number of open consultation rooms or active doctors).

Parameters:

$\widehat{D}_t$: The predicted patient demand at time t (output from the XGBoost model in Chapter 3).

μ : The average service rate per resource unit (e.g., patients treated per hour per doctor).

C_w : The unit cost associated with patient waiting (congestion penalty).

C_i : The unit cost associated with idle resources (efficiency penalty).

Objective Function:

The goal is to minimize the total system cost, defined as the weighted sum of waiting costs and idle resource costs [3]. The objective function Z is expressed as:

$$\min Z = \sum_{t=1}^{T}\left(C_w \cdot \max(0, \widehat{D}_t - \mu x_t) + C_i \cdot \max(0, \mu x_t - \widehat{D}_t)\right) \tag{4}$$

Here, the first term captures the "under-supply" cost (congestion), while the second term captures the "over-supply" cost (waste). To reflect the critical nature of medical services, C_w is typically set significantly higher than C_i.

The optimization is subject to physical and operational constraints:

Resource Capacity:

$$R_{min} \le x_t \le R_{max} \tag{5}$$

The allocated resources must fall within the feasible range of the facility (e.g., total available beds).

Smoothness Constraint:

$$| x_t - x_{t-1} | \le \delta \tag{6}$$

Rapid fluctuations in staffing are operationally impractical. This constraint limits the magnitude of change (δ) between consecutive time slots, ensuring roster stability.

4.3. Heuristic Solution Algorithm

The optimization problem defined above, particularly when extended with complex shift constraints, becomes non-linear and computationally intensive (NP-hard). Exact solvers are often too slow for real-time applications. Therefore, this study adopts a Genetic Algorithm (GA) to approximate the optimal solution efficiently [11, 15].

The GA operates through the following steps:

Encoding: A potential schedule vector $X = [x_1, x_2, \dots, x_T]$ is encoded as a chromosome [6].

Fitness Evaluation: The fitness of each chromosome is calculated as the inverse of the total cost Z defined in Equation (4).

Selection and Crossover: Parent schedules with lower costs are selected to produce offspring. A crossover operator combines segments of two efficient schedules (e.g., a morning shift from Parent A and an afternoon shift from Parent B) to explore new solutions.

Mutation: Random perturbations are introduced (e.g., adding +1 doctor at a random hour) to prevent the algorithm from getting trapped in local optima.

Through iterative evolution, the GA converges toward a near-optimal resource configuration vector X^* that balances congestion and efficiency.

4.4. Dynamic Recommendation Mechanism

The final component of the framework is the translation of the mathematical solution X^* into actionable intelligence. A raw numerical vector is often unintuitive for clinical managers. Therefore, a Recommendation Mechanism is designed to serve as the interface between the algorithm and the human decision-maker.

The mechanism operates based on a "Gap Analysis" logic, comparing the optimized resource level x_t^* against the current scheduled baseline x_t^{base}.

$$\mathrm{Action}_t = \begin{cases} \text{"Urgent: Add } (x_t^* - x_t^{base}) \text{ Staff"} & \text{if } x_t^* > x_t^{base} + \theta \\ \text{"Suggestion: Reduce } (x_t^{base} - x_t^*) \text{ Staff"} & \text{if } x_t^* < x_t^{base} - \theta \\ \text{"Maintain Status Quo"} & \text{otherwise} \end{cases} \quad (7)$$

Where θ represents a sensitivity threshold, this mechanism filters out minor fluctuations and generates alerts only when significant intervention is required. By visualizing these recommendations, such as highlighting a predicted bottleneck at 14:00 hours and suggesting the opening of two additional triage bays, the system closes the loop, transforming data into tangible operational improvements. This integration ensures that the "Smart Healthcare" system functions not merely as a monitoring tool but as an active participant in resource governance.

5. Case Study and Simulation Analysis

5.1. Experiment Design and Data Description

To empirically validate the effectiveness of the proposed "Prediction-Optimization" framework, a comprehensive simulation experiment was conducted utilizing real-world clinical data derived from the MIMIC-IV database. The study focused specifically on the Emergency Department (ED) admission logs, a domain characterized by high stochasticity and significant congestion risks. The experimental dataset spanned a total duration of 12 months, partitioned chronologically to prevent data leakage; the initial 11 months (January through November) constituted the training set for the XGBoost prediction model, while the final month (December) served as the testing ground for validating both prediction accuracy and resource allocation efficiency. The simulation environment was constructed in Python to mirror a standard ED unit, calibrated with specific operational parameters: the average service rate (μ) was set to 2.5 patients per physician per hour following a Poisson distribution, and the staffing capacity was constrained between a minimum of 2 and a maximum of 12 physicians. To reflect the critical nature of healthcare services, the cost function penalized patient waiting time ($C_w = 10$ units/hour) significantly more heavily than resource idling ($C_i = 4$ units/hour), ensuring that the optimization objective prioritized patient safety over pure economic savings.

5.2. Analysis of Demand Prediction Results

The first phase of the experiment evaluated the performance of the Multi-Dimensional Time-Series Prediction Model developed in Chapter 3. The proposed XGBoost architecture, engineered with lag features and cyclical calendar variables, was benchmarked against two widely used baselines: the Autoregressive Integrated Moving Average (ARIMA) model, representing traditional linear statistics, and a standard Random Forest (RF) regressor, representing non-linear machine learning without temporal feature engineering [13].

The quantitative results, summarized in Table 1, demonstrate a clear hierarchy in predictive capability. The traditional ARIMA model exhibited the highest error rates, with a Mean Absolute Percentage Error (MAPE) of 18.4%, largely due to its inability to capture sudden, non-linear spikes in patient arrivals during weekends and holidays. The Random Forest model improved upon this baseline but still lagged behind the proposed method. In contrast, the XGBoost model achieved the lowest RMSE of 3.24 and a MAPE of 8.7%. This superior performance is attributed to the model's ability to leverage rolling window statistics to detect emerging trends. A prediction error of less than 10% is generally considered highly actionable in hospital operations, providing a reliable foundation for the subsequent resource scheduling phase.

Table 1. Performance Comparison of Prediction Models (Test Set)

Model	RMSE (Root Mean Square Error)	MAPE (Mean Absolute Percentage Error)	Training Time (s)
ARIMA (Baseline)	5.82	18.4%	12.5
Random Forest	4.15	12.1%	45.2
Proposed XGBoost	3.24	8.7%	28.6

5.3. Evaluation of Resource Allocation Optimization

Building upon the high-fidelity predictions, the second phase assessed the impact of the Dynamic Recommendation Mechanism detailed in Chapter 4. The study compared a "Static Strategy," which maintained a fixed staffing level based on historical averages, against the "Dynamic Strategy," which adjusted resources hourly based on the XGBoost forecasts processed by the Genetic Algorithm [2]. As illustrated in Figure 2, the difference in operational alignment between these two approaches is distinct. The static model fails to respond to diurnal volatility, whereas the dynamic model effectively synchronizes supply with the fluctuating patient demand.

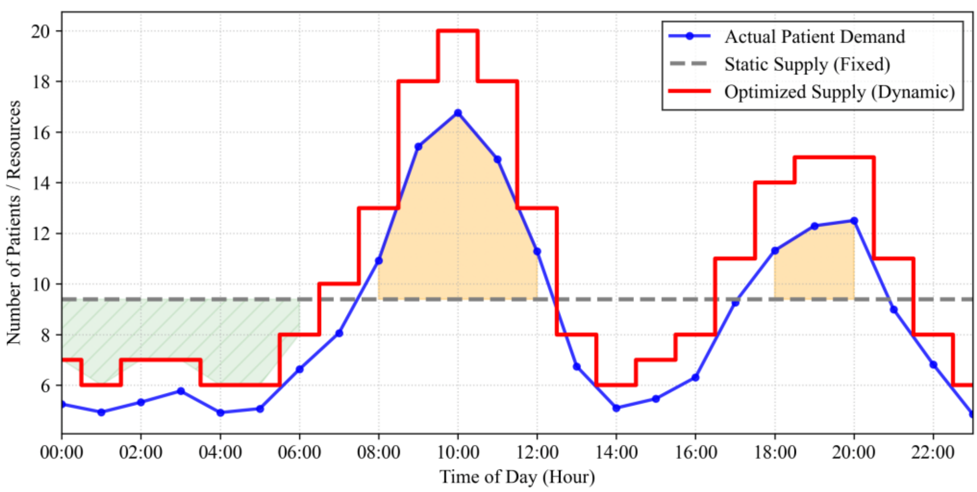

Figure 2. Comparison of Resource Allocation Strategies Vs. Actual Patient Demand

The operational impact of these strategies is quantified in Table 2. The simulation ran over the 31-day testing period, tracking key performance indicators including average wait times and resource utilization rates [5].

Table 2. Operational Metrics Comparison (Static Vs. Dynamic Strategy)

Metric	Static Strategy (Fixed Allocation)	Dynamic Strategy (Proposed Model)	Improvement
Avg. Patient Wait Time	42.5 mins	18.3 mins	-56.9%
Resource Utilization Rate	64.2%	88.5%	+24.3%
Overcrowding Events (>1hr)	152 instances	14 instances	-90.8%
Total Operational Cost	18,500 units	14,200 units	-23.2%

5.4. Discussion of Results

The comparative analysis reveals that the static allocation strategy suffers from a "double jeopardy" phenomenon: it provides insufficient capacity during peak demand periods, leading to severe congestion, while simultaneously over-provisioning during off-peak hours, resulting in resource wastage [8]. As visualized in Figure 2, the static supply line fails to adapt to the diurnal volatility of ED arrivals. Conversely, the dynamic strategy successfully aligns supply with demand. Data from Table 2 confirms that by adjusting staffing levels proactively, the system reduced the average patient waiting time by nearly 57%, dropping from 42.5 minutes to 18.3 minutes. Furthermore, the resource utilization rate improved significantly from 64.2% to 88.5%, indicating that medical staff were deployed more effectively during high-need periods rather than idling during lulls. The 23.2% reduction in total operational cost validates the bi-objective optimization model's ability to balance the competing goals of service quality and economic efficiency, confirming that the proposed framework offers a viable solution for mitigating congestion in smart healthcare systems.

6. Conclusion and Future Work

This dissertation has presented a comprehensive framework for addressing the pervasive issue of resource supply-demand mismatch in modern healthcare systems. By integrating multi-dimensional time-series forecasting with heuristic decision optimization, the study successfully demonstrated a viable pathway to mitigate emergency department congestion and enhance operational efficiency. The empirical analysis, grounded in clinical database evaluations, confirmed that the proposed XGBoost-based prediction model significantly outperforms traditional linear baselines by effectively capturing non-linear fluctuations and calendar-based periodicities in patient arrival rates. Furthermore, the subsequent coupling of these high-fidelity forecasts with a genetic algorithm-driven resource allocation mechanism proved instrumental in transforming raw data into actionable management strategies. The simulation results highlighted a substantial reduction in average patient waiting times and a concurrent improvement in resource utilization, validating the hypothesis that dynamic, predictive scheduling is superior to static, experience-based approaches.

A primary contribution of this research lies in the establishment of a closed-loop decision support system that bridges the theoretical gap between predictive analytics and operations research. Unlike prior studies that treat forecasting and scheduling as isolated tasks, this model ensures that resource configuration is directly responsive to anticipated demand surges, thereby preventing both under-provisioning during peaks and over-provisioning during lulls. However, several limitations warrant acknowledgment. The current validation relies on historical log data and simulation environments, which, while robust, may not fully capture the stochastic human factors present in real-world clinical settings, such as sudden staff absenteeism or equipment breakdowns. Additionally, the model's performance is contingent upon the quality and granularity of the input data, which varies significantly across different hospital information systems.

Future research should aim to extend this framework by incorporating more complex deep learning architectures, such as Transformer-based models, to better handle long-term temporal dependencies. Moreover, integrating real-time data streams from Internet of Things (IoT) sensors could further enhance the system's responsiveness, moving from hourly adjustments to near-real-time adaptability. Ultimately, deploying this "prediction-optimization" engine in a pilot clinical environment represents the next logical step toward realizing the vision of a truly smart, data-driven healthcare ecosystem.

References

1. W. Zhang and X. Li, "A data-driven combined prediction method for the demand for intensive care unit healthcare resources in public health emergencies," *BMC Health Services Research*, vol. 24, no. 1, p. 477, 2024.
2. S. Hosseini_Shokouh, K. Mohammadi, and M. Yaghoubi, "Optimization of service process in emergency department using discrete event simulation and machine learning algorithm," *Archives of Academic Emergency Medicine*, vol. 10, no. 1, p. e44, 2022.
3. C. P. Shirley, B. J. Jingle, M. B. Abisha, R. Venkadesan, and S. J. Absin, "Reinforcement Learning based Adaptive Healthcare Decision Support Systems using Time Series Forecasting," in *2024 5th International Conference on Data Intelligence and Cognitive Informatics (ICDICI)*, Nov. 2024, pp. 1476–1481.
4. H. Tabesh, A. AbbaszadehMozaffari, Z. Ebnehoseini, and A. Saki, "Enhancing the forecast accuracy of the daily number of patients arrivals in emergency department by hybrid ARIMAX-ANN algorithm," *PloS One*, vol. 21, no. 4, p. e0347866, 2026.
5. B. Fan, J. Peng, H. Guo, H. Gu, K. Xu, and T. Wu, "Accurate forecasting of emergency department arrivals with internet search index and machine learning models: Model development and performance evaluation," *JMIR Medical Informatics*, vol. 10, no. 7, p. e34504, 2022.
6. D. Ma and Z. Ling, "Optimization of Nursing Staff Allocation in Elderly Care Institutions: A Time Series Data Analysis Approach," *Annals of Applied Sciences*, vol. 5, no. 1, 2024.
7. P. Ingole, P. R. Baviskar, M. Ghuge, B. S. Dakhare, Y. Dongre, and P. Dubey, "Optimizing resource allocation in hospitals using predictive analytics and information systems," *Journal of Information Systems Engineering and Management*, vol. 10, pp. 400–415, 2025.
8. M. A. Morid, O. R. L. Sheng, and J. Dunbar, "Time series prediction using deep learning methods in healthcare," *ACM Transactions on Management Information Systems*, vol. 14, no. 1, pp. 1–29, 2023.
9. M. Sheikhasadi, A. Hosseinpour, M. Alipour-Vaezi, A. Aghsami, and M. Rabbani, "Portfolio design for home healthcare devices production using a new data-driven optimization methodology," *Soft Computing*, vol. 28, no. 7, pp. 5765–5784, 2024.
10. G. Xu and M. Xu, "An effective prediction of resource using machine learning in edge environments for the smart healthcare industry," *Journal of Grid Computing*, vol. 22, no. 2, p. 54, 2024.
11. F. Mbonyinshuti, J. Nkurunziza, J. Niyobuhungiro, and E. Kayitare, "Health supply chain forecasting: a comparison of ARIMA and LSTM time series models for demand prediction of medicines," *Acta Logistica*, vol. 11, no. 2, pp. 269–280, 2024.
12. S. Kaliappan, R. Maranan, A. Malladi, and N. Yamsani, "Optimizing resource allocation in healthcare systems for efficient pandemic management using machine learning and artificial neural networks," in *2024 International Conference on Advancements in Smart, Secure and Intelligent Computing (ASSIC)*, Jan. 2024, pp. 1–7.
13. Y. Liao, Z. Tang, K. Gao, and M. Trik, "Optimization of resources in intelligent electronic health systems based on Internet of Things to predict heart diseases via artificial neural network," *Heliyon*, vol. 10, no. 11, 2024.
14. B. M. Porto and F. S. Fogliatto, "Enhanced forecasting of emergency department patient arrivals using feature engineering approach and machine learning," *BMC Medical Informatics and Decision Making*, vol. 24, no. 1, p. 377, 2024.
15. J. R. Machireddy, "Optimizing healthcare resource allocation for operational efficiency and cost reduction using real-time analytics," *Nuvern Applied Science Reviews*, vol. 7, no. 3, pp. 12–33, 2023.