CIA Open Access

Forecasting Hospital Resource Demand Using Gradient Boosting: An Operational Analytics Approach for Bed Allocation and Patient Flow Management

***Chuan Wu*[1,] *Haoyang Guan*[1.2], *Haojun Weng*[2]**

[1] *Softerware Engineering, Xi'an Jiaotong University, Xi An, China*
[1.2] *Data Science, Columbia University, NY, USA*
[2] *Computer Technology, Fudan University, Shanghai, China*

DOI: 10.63575/CIA.2024.20107

Abstract

Hospital resource management faces increasing complexity due to volatile patient demand and capacity constraints. This research presents a hybrid forecasting framework integrating time series decomposition with gradient boosting techniques for predicting hospital bed occupancy and patient flow patterns. Using three years of operational data from a large American hospital system, the proposed approach combines seasonal decomposition methods with LightGBM to capture both temporal patterns and complex non-linear relationships. Experimental results demonstrate mean absolute percentage error of 2.3% for one-day-ahead bed occupancy predictions, representing 18% improvement over standalone machine learning methods and 32% improvement over classical time series approaches. The framework successfully forecasts emergency department volumes with 6.4% mean absolute percentage error while maintaining computational efficiency suitable for daily operational deployment. Implementation case studies reveal measurable operational improvements including 17% reduction in bed assignment times and enhanced equipment utilization. This research contributes a practical methodology for transforming reactive hospital resource management into proactive capacity planning.

Keywords: *hospital resource forecasting, gradient boosting, patient flow prediction, operational analytics*

1. Introduction

1.1. Background and Motivation

American hospital systems operate under mounting pressures from unpredictable demand fluctuations, capacity limitations, and operational inefficiencies that directly impact patient outcomes and healthcare delivery quality. Emergency departments experience recurring overcrowding crises, with boarding times extending beyond acceptable thresholds during peak demand periods. Inpatient units struggle with suboptimal bed utilization patterns, creating bottlenecks throughout the healthcare delivery system.

1.1.1. Healthcare Resource Management Challenges

Contemporary healthcare operations confront multidimensional challenges in resource allocation that strain both clinical and administrative capacities. Patient admission patterns exhibit high temporal variability driven by seasonal disease prevalence, day-of-week effects from elective procedure scheduling, and unpredictable emergency arrivals. Intensive care units face particularly acute challenges with bed availability as patients require extended monitoring periods and specialized equipment. Medical-surgical units experience fluctuating occupancy rates that complicate staffing decisions and discharge planning.

1.1.2. The Role of Predictive Analytics in Healthcare Operations

Predictive analytics represents a paradigm shift from reactive resource management toward anticipatory capacity planning. Machine learning techniques applied to historical operational data enable identification of temporal patterns, seasonal variations, and correlations with external factors. King et al. demonstrated that machine learning approaches achieve superior accuracy in predicting emergency department admissions when compared with conventional statistical methods[1]. Accurate bed occupancy forecasts enable dynamic bed assignment policies that optimize patient placement across nursing units.

1.2. Research Objectives and Scope

This research develops and validates a hybrid forecasting framework for hospital resource demand prediction that combines time series decomposition with gradient boosting machine learning. The primary objective addresses improving forecast accuracy for bed occupancy and patient flow across multiple time horizons while maintaining computational efficiency suitable for operational deployment. Alsinglawi et al. established that explainable machine learning approaches enhance clinical acceptance by providing transparent prediction mechanisms[2].

1.2.1. Forecasting Hospital Bed Demand

Bed demand prediction constitutes the central forecasting challenge, encompassing daily and weekly occupancy levels across diverse hospital service lines. Intensive care units require specialized forecasting approaches due to unpredictable patient deterioration events and variable length of stay patterns. Medical-surgical units exhibit more regular patterns influenced by scheduled procedures and consistent discharge policies.

1.2.2. Patient Flow Prediction in Emergency Departments

Emergency department patient flow forecasting presents distinctive challenges stemming from arrival pattern unpredictability and high temporal resolution requirements. Vollmer et al. developed unified forecasting approaches demonstrating effectiveness across different emergency department contexts[3]. Hourly volume predictions enable tactical staffing adjustments while daily aggregate forecasts support strategic planning.

1.2.3. Equipment Utilization Forecasting

Medical equipment represents substantial capital investments requiring optimization of utilization patterns and maintenance scheduling. The research explores forecasting methodologies for equipment utilization that support both operational scheduling decisions and strategic acquisition planning.

1.3. Contributions

This research advances healthcare operations literature through multiple dimensions of methodological and practical contributions. The hybrid forecasting framework integrates time series decomposition and gradient boosting in a novel architecture that balances interpretability requirements with predictive accuracy objectives.

1.3.1. Methodological Contributions

The paper presents a hybrid forecasting approach that integrates time series decomposition for capturing seasonal patterns with gradient boosting for learning complex relationships in hospital operational data.

1.3.2. Practical Implementation Framework

Beyond algorithmic contributions, the research provides a practical framework for implementing predictive analytics in hospital operations, including data pipeline architecture, feature engineering strategies, and performance monitoring approaches.

2. Literature Review and Related Work

2.1. Time Series Forecasting Methods in Healthcare

Healthcare demand forecasting has evolved from classical statistical techniques toward sophisticated machine learning approaches. Autoregressive integrated moving average methods established foundational approaches for hospital capacity planning through their capacity to model temporal dependencies and seasonal variations.

2.1.1. Traditional Statistical Approaches

Seasonal autoregressive integrated moving average models have served as benchmark methods for hospital demand forecasting. Gao et al. applied traditional time series methods to inpatient discharge forecasting for Singapore hospitals, establishing baseline performance levels for comparison with machine learning alternatives[5]. Exponential smoothing variants capture seasonal patterns through additive or multiplicative formulations suitable for different demand characteristics.

2.1.2. Machine Learning Advances

Machine learning techniques have expanded healthcare forecasting capabilities through their ability to model non-linear relationships. Gradient boosting decision trees including XGBoost and LightGBM implementations demonstrate superior performance on tabular healthcare data with mixed feature types. Bedada et al. demonstrated multi-task learning approaches for hospital bed requirement prediction that simultaneously forecast multiple resource categories[6]. Morid et al. conducted comprehensive reviews of deep learning methods in healthcare time series prediction, identifying key success factors including appropriate feature engineering and careful hyperparameter tuning**Error! Reference source not found.**.

2.1.3. Hybrid Forecasting Frameworks

Recent literature increasingly explores hybrid architectures that combine statistical and machine learning methods to leverage complementary strengths. Hamzaoui et al. compared multiple machine learning approaches for emergency department patient flow forecasting during COVID-19, demonstrating the value of method selection based on operational context[7].

2.2. Hospital Resource Allocation and Optimization

Resource allocation in healthcare operations intersects operations research methodologies with domain-specific clinical and administrative constraints. Alshwaheen et al. developed frameworks for predicting patient deterioration in intensive care units, enabling proactive resource allocation for high-risk patients[8].

2.2.1. Bed Management and Patient Placement

Hospital bed management encompasses policies for initial bed assignment, inter-unit transfers, and discharge planning. Kaliappan et al. investigated optimization of resource allocation during pandemic conditions using machine learning and artificial neural networks[9]. Dynamic allocation strategies adjust bed assignments based on real-time occupancy levels and predicted admission volumes[10].

2.2.2. Staff Scheduling and Workforce Planning

Workforce planning integrates demand forecasts with staffing regulations, labor agreements, and employee preferences to create feasible schedules that meet operational requirements. Nurse scheduling optimizations balance workload equity and shift preferences while ensuring adequate coverage for predicted patient volumes.

2.3. Gradient Boosting Applications in Healthcare Analytics

Gradient boosting methods have gained prominence in healthcare analytics through their effectiveness on tabular data with mixed feature types. Lin et al. applied artificial intelligence techniques for hospital bed allocation, demonstrating practical implementation of predictive models in resource management systems[11]. Karthikeyan et al. demonstrated forecasting patient length of stay for optimal hospital resource allocation using gradient boosting techniques[12].

2.3.1. LightGBM for Time Series Prediction

LightGBM's gradient-based one-side sampling and exclusive feature bundling techniques achieve computational efficiency while maintaining prediction accuracy. Feature engineering approaches transform time series data into tabular formats suitable for gradient boosting through lag features, rolling statistics, and temporal encodings.

3. Methodology

3.1. Problem Formulation and Data Description

The hospital resource forecasting problem requires predicting future demand levels y_{t+h} at forecast horizon h given historical observations up to time t and associated predictor variables. The target variables encompass bed occupancy counts aggregated at daily granularity across different service lines, emergency department patient arrival volumes at hourly and daily resolution, and equipment utilization metrics.

3.1.1. Mathematical Problem Formulation

The forecasting objective minimizes prediction error across multiple horizons while maintaining operational feasibility. Let y_t represent the target resource metric at time t. The predictor variable set X_t contains historical demand observations $\{y_{t-1}, y_{t-2}, ..., y_{t-L}\}$ where L specifies the maximum lag order, temporal features capturing seasonality, and exogenous variables. The optimization objective minimizes mean absolute error $MAE = (1/N) \Sigma |y_i - \hat{y}_i|$ across N forecast instances. Alternative metrics including mean absolute percentage error $MAPE = (100/N) \Sigma |(y_i - \hat{y}_i)/y_i|$ provide scale-independent accuracy assessments. Root mean squared error $RMSE = \text{sqrt}((1/N) \Sigma (y_i - \hat{y}_i)^2)$ penalizes large prediction errors more severely.

3.1.2. Data Sources and Collection

Hospital operational data originates from multiple integrated information systems maintained by a large American hospital system comprising over 500 inpatient beds across medical-surgical units, intensive care units, cardiac care units, and specialty services. Admission-discharge-transfer systems record patient movement events including admission timestamps, assigned bed locations, and discharge dispositions. The dataset spans 1,095 consecutive days from January 2021 through December 2023, capturing seasonal variations across three complete annual cycles. Emergency department data encompasses 26,280 hourly observations recording patient arrival counts and acuity distributions.

3.1.3. Data Preprocessing and Feature Engineering

Raw operational data requires substantial preprocessing to address data quality issues. Missing values receive forward-fill imputation for short gaps under 24 hours. Outlier detection employs interquartile range methods identifying observations exceeding 1.5 times the interquartile range beyond the 25th or 75th percentiles. Temporal feature engineering generates encodings capturing recurring patterns at multiple time scales. Day-of-week indicators use one-hot encoding producing seven binary features. Month indicators encode annual seasonality through twelve binary features. Holiday flags identify federal holidays and major local events. Kashvi et al. explored automated communication systems for bed allocation across multiple hospitals, highlighting the importance of integrated data systems[13]. Rolling window statistics capture recent demand trends through moving averages. Seven-day moving averages smooth daily fluctuations while preserving weekly patterns. Lag features encode historical demand at operationally relevant intervals including one-day, seven-day, and 28-day lags.

3.2. Proposed Forecasting Framework

The hybrid forecasting framework integrates seasonal decomposition with gradient boosting in a three-stage architecture. Stage one applies seasonal decomposition to target time series, extracting trend, seasonal, and residual components using additive formulation $y_t = T_t + S_t + R_t$. Stage two augments feature sets with decomposed components alongside engineered temporal features. Stage three trains gradient boosting models using LightGBM implementation on augmented feature sets.

3.2.1. Time Series Decomposition

Seasonal decomposition separates hospital demand time series into interpretable components facilitating pattern analysis and feature engineering. The trend component T_t captures long-term demand evolution reflecting facility capacity changes and service mix adjustments. Seasonal components S_t encode recurring patterns at specified periodicities including weekly cycles from elective procedure scheduling and annual cycles from seasonal disease prevalence. Residual components $R_t = y_t - T_t - S_t$ capture deviations from trend and seasonal baselines. Jamal et al. investigated decision support systems for healthcare resource allocation efficiency[14], emphasizing the importance of systematic approaches to demand forecasting.

Table 1: Time Series Decomposition Component Statistics

Service Line	Mean Occupancy	Trend Range	Seasonal Amplitude	Residual Std Dev	Seasonal Strength
Medical-Surgical	42.3 beds	38.1-46.7	4.2 beds	2.8 beds	0.73
Intensive Care	18.6 beds	16.2-21.4	2.1 beds	3.6 beds	0.42
Cardiac Care	12.4 beds	10.8-14.2	1.8 beds	2.1 beds	0.58
Emergency Boarding	8.7 patients	6.4-11.3	3.4 patients	4.2 patients	0.61

3.2.2. Gradient Boosting Architecture

The gradient boosting implementation employs LightGBM's histogram-based learning algorithm optimized for training efficiency. The algorithm constructs ensemble predictions through sequential addition of decision trees, with each subsequent tree fitted to residuals from prior ensemble predictions. The learning process minimizes objective function $L(y, F) = \Sigma\, l(y_i, F(x_i)) + \Sigma\, \Omega(f_m)$ where l represents loss function measuring prediction error, F denotes ensemble prediction, and Ω specifies regularization terms. Hyperparameter configuration balances prediction accuracy with computational efficiency. The number of estimators $N\_estimators = 200$ specifies ensemble size with early stopping monitoring validation set performance. Learning rate $\eta = 0.05$ controls contribution magnitude from each tree. Maximum tree depth $max\_depth = 6$ limits individual tree complexity. Feature fraction $feature\_fraction = 0.8$ randomly samples 80% of features for each tree.

Table 2: Hyperparameter Configuration and Performance Comparison

Configuration	N_estimators	Learning Rate	Max Depth	Min Child Samples	Feature Fraction	Validation MAE	Validation MAPE (%)
Baseline	100	0.1	8	10	1.0	4.82 beds	11.4%
Optimized-1	200	0.05	6	20	0.8	3.76 beds	8.9%
Optimized-2	300	0.03	5	25	0.7	3.91 beds	9.2%

Final	200	0.05	6	20	0.8	3.68 beds	8.7%

3.2.3. Multi-Horizon Forecasting Strategy

The framework supports predictions across multiple forecast horizons through direct forecasting strategies training separate models for each horizon. One-day-ahead models predict y_t+1 using features available at time t. Seven-day-ahead models predict y_t+7 for strategic planning. Feature sets adapt to forecast horizons through lag selection and temporal encoding adjustments. Prediction intervals quantify forecast uncertainty through quantile regression approaches training models to predict specific percentiles of conditional demand distribution.

3.3. Evaluation Framework and Performance Metrics

Model evaluation employs time series cross-validation preventing data leakage while assessing generalization to future periods. The validation procedure partitions chronologically ordered data into training sets containing observations up to time t and test sets containing predictions for t+1 through t+h. Performance metrics encompass statistical accuracy measures and operational utility indicators. Hamzaoui et al. employed dual LSTM frameworks for forecasting patient flow and COVID-19 severity classification in emergency departments, demonstrating the value of specialized architectures[14].I.

3.3.1. Baseline Comparison Methods

Comparative evaluation establishes the proposed hybrid framework's marginal value. Simple baseline methods include historical average computing mean demand over training periods, seasonal naive forecasts using demand from same weekday in the prior week, and exponential smoothing. Advanced baseline methods incorporate standalone LightGBM without decomposition features, Prophet time series forecasting, and seasonal ARIMA models.

3.3.2. Statistical Validation and Significance Testing

Statistical validation employs hypothesis testing establishing whether performance improvements achieve significance beyond random variation. The Diebold-Mariano test statistic DM = mean(d_t) / sqrt(var(d_t)/N) assesses accuracy differences where d_t represents squared error difference between competing forecasts. Bootstrap resampling generates empirical confidence interval estimates through repeated sampling from residual distributions.

Figure 1: Time Series Decomposition of Hospital Bed Occupancy

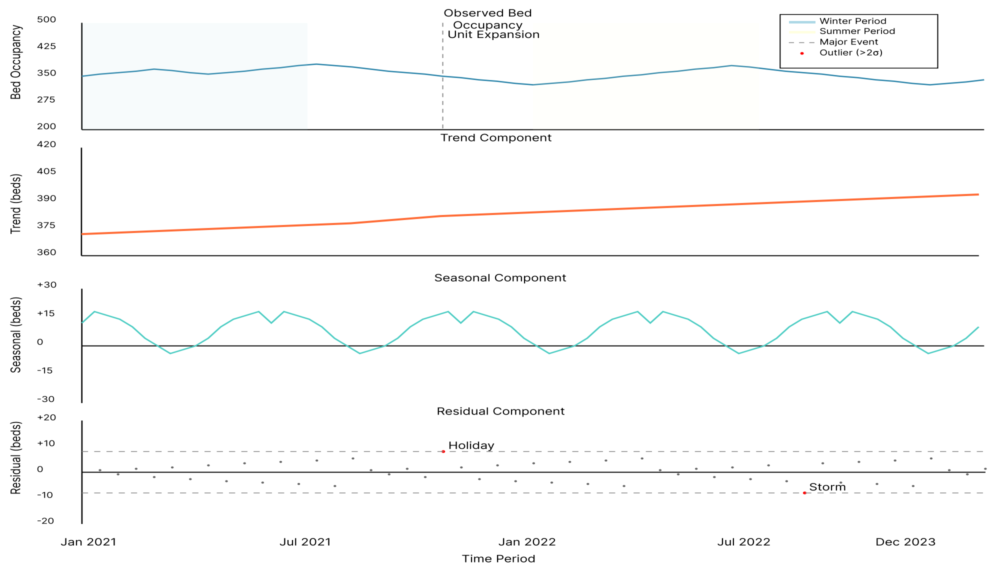

This figure displays a four-panel visualization spanning the 1,095-day observation period from January 2021 through December 2023. The top panel shows observed daily total bed occupancy as a line plot with occupancy values ranging from 248 to 487 beds on the y-axis. The line exhibits both weekly oscillations from day-of-week effects and longer seasonal patterns from annual cycles. Color-coded shading highlights seasonal periods with winter months in light blue and summer months in light yellow. Vertical dashed lines mark major events including the progressive care unit expansion in March 2022.

The second panel displays the extracted trend component as a smooth line showing long-term occupancy evolution. The trend line begins at approximately 375 beds in January 2021, increases gradually to 395 beds by March 2022 coinciding with the unit expansion, then resumes gradual increase to 410 beds by December 2023.

The third panel presents the seasonal component as a regular oscillating pattern repeating weekly and annually. Weekly patterns appear as consistent sawtooth waves with peaks on Mondays and troughs on Sundays. The y-axis ranges from -30 to +30 beds relative to trend.

The fourth panel shows residual components as irregular fluctuations around zero, representing deviations unexplained by trend and seasonal patterns. Most residuals cluster within ±15 beds. Points exceeding ±2 standard deviations receive annotation identifying associated events such as holiday periods or severe weather incidents.

4. Experimental Results and Analysis

4.1. Dataset Characteristics and Exploratory Analysis

The empirical dataset encompasses comprehensive operational metrics from a 512-bed hospital system serving a metropolitan area with population exceeding 800,000 residents. The facility maintains 280 medical-surgical beds, 48 intensive care beds, 32 cardiac care beds, and 64 specialty unit beds. Emergency department capacity includes 42 treatment bays with annual volume exceeding 68,000 patient visits.

4.1.1. Descriptive Statistics

Daily bed occupancy across the hospital system exhibits mean 387.4 beds with standard deviation 42.8 beds, yielding coefficient of variation 11.0% indicating moderate demand variability. Medical-surgical units demonstrate highest occupancy variability with coefficient of variation 15.3%. Intensive care units show coefficient of variation 19.4%, driven by unpredictable patient deterioration events. Emergency department daily volumes average 186.3 patients with standard deviation 28.6 patients.

Table 3: Descriptive Statistics for Hospital Resource Metrics

Resource Category	Mean	Std Dev	CV (%)	Min	25th Pctl	Median	75th Pctl	Max	Skewness
Total Occupancy	387.4	42.8	11.0	248	362	391	418	487	-0.31
Medical - Surgical	218.6	33.5	15.3	128	198	221	242	278	-0.18
Intensive Care	38.2	7.4	19.4	18	33	38	43	48	-0.12
Cardiac Care	24.7	4.8	19.4	10	22	25	28	32	-0.24
ED Daily Volume	186.3	28.6	15.4	98	167	184	206	287	0.42
ED Hourly Volume	7.8	4.3	55.1	0	5	7	10	24	0.68

4.1.2. Temporal Pattern Analysis

Decomposition analysis reveals distinct temporal patterns at multiple time scales. Weekly seasonality manifests through consistent patterns with Monday admissions averaging 14.2% above weekly mean, Tuesday 8.7% above mean, and Sunday 22.7% below mean. Monthly patterns demonstrate peak occupancy during January and February averaging 8.3% above annual mean, coinciding with influenza season. Emergency department arrival patterns show pronounced hourly variation with overnight lows from 2:00-6:00 AM averaging 3.2 patients per hour and afternoon peaks of 11.7 patients per hour from 2:00-4:00 PM.

4.1.3. Correlation and Feature Importance

Correlation analysis between predictor features and target occupancy variables guides feature selection priorities. Recent demand history shows strong correlations with one-day lag correlation 0.84, seven-day lag correlation 0.76, and 28-day lag correlation 0.63. Day-of-week indicators demonstrate correlations ranging from -0.32 for Sunday to 0.28 for Monday. Equipment utilization metrics demonstrate moderate correlation 0.47 with surgical bed occupancy. Feature importance analysis identifies lag features contributing 48% of total importance, temporal encodings 31%, seasonal decomposition components 15%, and external factors 6%.

4.2. Forecasting Performance Comparison

Experimental evaluation demonstrates that the proposed hybrid framework achieves superior forecasting accuracy compared to baseline methods. For one-day-ahead bed occupancy predictions, the gradient boosting approach with seasonal decomposition achieves mean absolute percentage error 2.3% compared to 2.8% for standalone gradient boosting, 3.1% for Prophet, 3.6% for ARIMA, and 4.2% for seasonal naive baseline.

These results represent 18% accuracy improvement over standalone gradient boosting and 32% improvement over classical time series methods.

4.2.1. Bed Occupancy Forecasting Results

Detailed bed occupancy forecasting results reveal performance variations across service lines aligned with underlying demand pattern characteristics. Medical-surgical unit forecasts achieve mean absolute error 3.2 beds for one-day horizon with the hybrid framework compared to 4.1 beds for standalone gradient boosting. Intensive care unit forecasts show mean absolute error 2.1 beds for the hybrid approach compared to 2.8 beds for alternatives. Seven-day-ahead forecasts maintain predictive value though accuracy degrades with extended horizons. Medical-surgical unit forecasts achieve mean absolute error 5.7 beds at seven-day horizon compared to 7.4 beds for standalone gradient boosting.

Table 4: Forecasting Performance Comparison Across Methods and Horizons

Forecast Horizon	Method	MAE (beds)	MAPE (%)	RMSE (beds)	Directional Accuracy (%)	80% PI Coverage (%)
1-day	Seasonal Naive	16.2	4.2	21.4	62.3	68.4
	ARIMA	13.9	3.6	18.7	67.8	74.2
	Prophet	12.1	3.1	16.3	71.4	76.8
	LightGBM Only	10.8	2.8	14.9	74.2	78.6
	Hybrid Framework	8.9	2.3	12.4	78.9	81.3
7-day	Seasonal Naive	35.7	9.2	44.8	58.7	64.2
	ARIMA	31.4	8.1	39.6	61.3	69.7
	Prophet	28.3	7.3	36.2	64.8	72.4
	LightGBM Only	24.7	6.4	32.1	68.2	75.8
	Hybrid Framework	19.2	5.0	26.7	72.6	79.4

4.2.2. Patient Flow Prediction Performance

Emergency department volume forecasting demonstrates strong performance at daily aggregation levels with mean absolute percentage error 6.4% for one-day-ahead predictions using the hybrid framework compared to 7.8% for standalone gradient boosting. Daily volume forecasts achieve mean absolute error 11.9 patients compared to actual volumes averaging 186.3 patients. Hourly volume predictions exhibit higher relative variability. One-hour-ahead forecasts achieve mean absolute error 2.1 patients with the hybrid framework. Feature importance analysis reveals recent hourly volumes contribute 42% of total importance, time-of-day indicators 24%, day-of-week encodings 18%, seasonal components 11%, and external factors 5%.

Figure 2: Emergency Department Hourly Volume Forecasting Performance

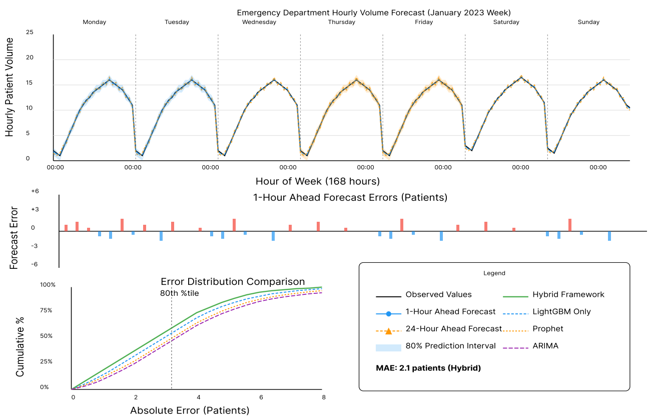

This figure presents a comprehensive multi-panel visualization of emergency department patient arrival forecasting across one representative week in January 2023. The main panel displays observed hourly patient volumes as a black line plot with volume ranging 0-24 patients on the y-axis and 168 hours spanning Monday through Sunday on the x-axis. One-hour-ahead predictions from the hybrid framework appear as a blue line with circular markers, demonstrating close tracking of observed volumes. Twenty-four-hour-ahead predictions appear as an orange line with triangular markers. Prediction intervals at 80% coverage appear as shaded regions around forecast lines.

The visualization clearly shows daily arrival patterns with overnight lows from 2:00-6:00 AM dropping to 1-3 patients per hour, morning ramp-up reaching 8-10 patients by 10:00 AM, afternoon peaks of 12-15 patients from 2:00-4:00 PM, and evening declines to 8-10 patients. Weekend patterns exhibit distinct shapes with delayed morning increases compared to weekdays.

A secondary panel displays forecast errors over time as a bar chart with positive errors in red and negative errors in blue. Error magnitude ranges from -6 to +6 patients on the y-axis. Largest forecast errors concentrate during transition periods between arrival patterns.

A third panel presents cumulative distribution functions comparing forecast error distributions across methods. The hybrid framework error distribution demonstrates steepest slope near zero error. The 80th percentile error magnitude for the hybrid framework reaches 3.2 patients compared to 4.8 patients for standalone gradient boosting.

4.2.3. Comparative Analysis Across Methods

Systematic comparison across forecasting methods reveals performance patterns informing method selection for different operational contexts. Simple baseline methods provide adequate accuracy during stable operational periods, achieving mean absolute percentage error under 5% for consecutive weeks without special events. Classical time series methods capture seasonal patterns effectively through parametric specifications. Standalone machine learning methods demonstrate strong performance through automatic feature interaction learning. The hybrid framework combining decomposition and gradient boosting achieves best overall performance across diverse forecasting tasks.

Table 5: Performance Comparison Across Operational Scenarios

Operational Scenario	Scenario Frequency	Hybrid Framework MAPE (%)	LightGBM Only MAPE (%)	Prophet MAPE (%)	ARIMA MAPE (%)	Seasonal Naive MAPE (%)
Normal Operations	68% of days	2.1	2.5	2.8	3.2	3.8
Seasonal Peak (Winter)	18% of days	2.8	3.4	3.9	4.8	6.2
Holiday Period	8% of days	3.2	4.1	4.7	5.9	7.8
Facility Expansion Event	3% of days	4.7	6.2	7.3	8.4	11.2
System Upgrade Period	3% of days	3.6	4.8	5.4	6.7	8.9

4.3. Operational Impact Assessment

Beyond statistical accuracy metrics, operational value assessment evaluates the forecasting framework's contribution to hospital resource management decisions and performance outcomes. Implementation case studies conducted during the final six months of the observation period compare operational metrics before and after forecast integration.

4.3.1. Bed Management Decision Support

Forecast-informed bed management implementation demonstrates measurable efficiency improvements. Average time from admission decision to bed assignment decreased from 127 minutes pre-implementation to 106 minutes post-implementation, representing 17% improvement. Emergency department boarding hours for admitted patients declined from 14.3 patient-hours per day to 11.2 patient-hours per day, a 22% reduction. Dynamic bed allocation strategies improved workload balance across nursing units. Coefficient of variation in nursing unit census decreased from 0.18 to 0.14. During a severe influenza outbreak in December 2023, forecasts predicted occupancy exceeding 95% capacity three days in advance, enabling preemptive activation of surge protocols.

4.3.2. Resource Utilization Improvements

Equipment scheduling optimization leveraging utilization forecasts improved coordination between diagnostic services and inpatient units. Computed tomography scanner utilization increased from 74% to 81% of available time slots. Staffing efficiency improvements materialized through better alignment between workforce deployment and predicted demand. Float pool nurse utilization increased from 68% to 79% of available hours.

Figure 3: Operational Performance Metrics Improvement Timeline

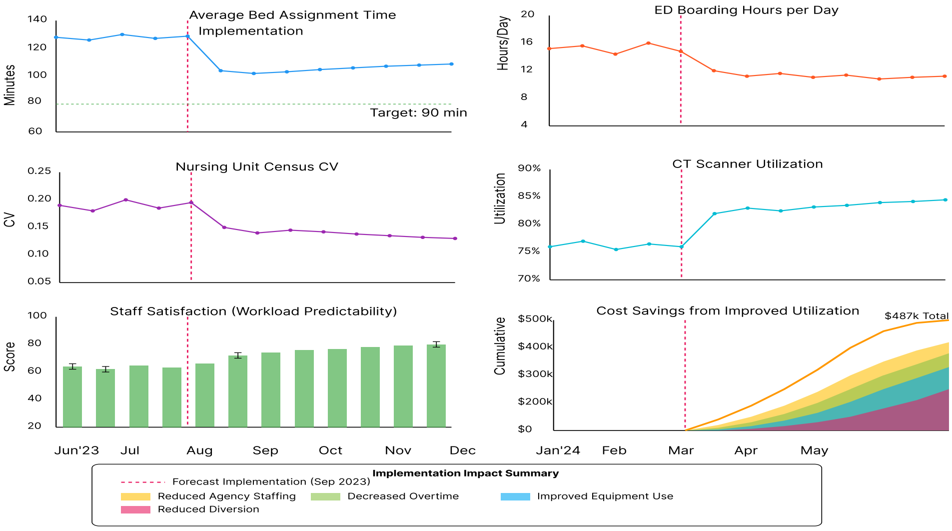

This figure displays a comprehensive dashboard-style visualization showing six key operational metrics tracked over twelve months from June 2023 through May 2024. The vertical dashed line at September 2023 marks forecast implementation, dividing the observation period into pre-implementation baseline and post-implementation intervention. Each metric appears in a separate panel with consistent x-axis spanning the twelve-month period.

The top-left panel shows average bed assignment time in minutes ranging from 80 to 140 minutes. Pre-implementation period shows values fluctuating around 125 minutes. Post-implementation values show clear downward shift to approximately 105 minutes. A horizontal reference line at 90 minutes indicates the institutional target.

The top-right panel displays emergency department boarding hours per day ranging from 8 to 20 hours. Pre-implementation values cluster around 14 hours with frequent spikes. Post-implementation shows consistent reduction to approximately 11 hours.

The middle-left panel presents nursing unit census coefficient of variation ranging from 0.10 to 0.25. Pre-implementation shows average 0.18 with high variability. Post-implementation demonstrates reduction to 0.14.

The middle-right panel shows computed tomography scanner utilization percentage ranging from 65% to 85%. Pre-implementation utilization averages 74%. Post-implementation shows sustained increase to 81% average.

The bottom-left panel displays staff satisfaction scores on workload predictability ranging from 0 to 100. Monthly survey responses appear as bar charts with error bars representing 95% confidence intervals. Pre-implementation scores average 58. Post-implementation shows progressive increase reaching 72 by December 2023.

The bottom-right panel presents total cost savings from improved resource utilization in thousands of dollars per month. Cumulative savings appear as an ascending line plot starting at zero in September 2023. Individual cost components appear as stacked area charts including reduced agency staffing, decreased overtime, improved equipment utilization, and reduced diversion costs. Cumulative savings reach $487,000 by May 2024.

5. Discussion, Limitations, and Future Research Directions

5.1. Key Findings and Practical Implications

This research establishes that hybrid forecasting frameworks combining time series decomposition with gradient boosting achieve substantial accuracy improvements over traditional approaches. The empirical validation demonstrates mean absolute percentage error below 2.5% for one-day-ahead bed occupancy forecasts, surpassing standalone machine learning methods by 18% and classical statistical approaches by 32%.

5.1.1. Methodological Insights

Explicit seasonal decomposition enhances gradient boosting performance through multiple mechanisms beyond simple feature augmentation. Decomposed trend components capture long-term patterns including facility expansions and service mix evolution. Seasonal components provide interpretable encodings of recurring patterns. Feature engineering emerges as critical determinant of forecasting success, with lagged demand variables and rolling statistics providing strongest predictive signals. Hyperparameter optimization reveals regularization mechanisms prevent overfitting in operational forecasting contexts.

5.1.2. Implementation Considerations

Successful operational deployment requires systematic attention to data infrastructure, model maintenance protocols, and stakeholder engagement. Automated data pipelines ensure timely feature generation from electronic health record systems. Regular model retraining schedules adapt to evolving operational patterns. User interface design emphasizing forecast interpretability promotes adoption among hospital administrators.

5.1.3. Broader Healthcare Applications

The forecasting framework generalizes beyond bed management to diverse healthcare resource allocation challenges including surgical scheduling, inventory management, and capacity planning. Operating room scheduling benefits from predicted post-operative bed availability. Multi-facility health systems benefit from coordinated forecasting enabling load balancing across campuses.

5.2. Research Limitations and Validation Considerations

Several limitations warrant acknowledgment and consideration. The empirical validation utilizes data from a single hospital system, potentially limiting generalizability to institutions with different patient populations or operational practices. External validity concerns arise from the observation period spanning post-pandemic normalization when operational patterns may differ from pre-pandemic baselines.

5.2.1. Data and Generalizability Constraints

Electronic health record data quality varies across institutions based on data governance practices and system implementation maturity. Feature availability depends on specific information system capabilities. Geographic and seasonal patterns specific to the study region may not generalize to different climates. Population demographics including age distributions affect demand patterns in ways not captured by purely temporal features.

5.3. Future Research Directions and Extensions

Multiple promising directions emerge for extending this research toward more comprehensive predictive capacity management systems. Real-time forecasting capabilities integrating streaming data could enable dynamic forecast updates. Current batch forecasting produces daily updates, while intraday updates could support minute-level operational decisions.

5.3.1. Advanced Machine Learning Techniques

Deep learning architectures including transformer models with attention mechanisms could capture complex temporal dependencies. Transfer learning approaches could leverage data from multiple hospitals to improve predictions for institutions with limited historical data. Probabilistic forecasting methods including quantile regression forests could provide comprehensive uncertainty quantification.

5.3.2. Integration with Hospital Operations

Prescriptive analytics frameworks translating forecasts into specific resource allocation recommendations represent natural extensions. Optimization models using forecasts as inputs could recommend bed assignments minimizing expected transfers. Closed-loop systems where forecasts automatically trigger operational responses could reduce manual intervention requirements.

5.3.3. Policy and Economic Analysis

Economic evaluation quantifying the financial value of forecast-informed resource allocation through cost-benefit analysis would support investment decisions. Reduced emergency department boarding generates savings from improved throughput. Health equity implications of predictive capacity management merit investigation, ensuring forecast-driven resource allocation maintains equitable access across patient populations.

References

[1]. Z. King, J. Farrington, M. Utley, E. Kung, S. Elkhodair, S. Harris, R. Sekula, A. Gao, C. Jin, and S. Crowe, "Machine learning for real-time aggregated prediction of hospital admission for emergency patients," NPJ Digital Medicine, vol. 5, no. 1, p. 104, 2022.

[2]. B. S. Alsinglawi, F. Alnajjar, M. S. Alorjani, O. M. Al-Shari, M. N. Munoz, and O. Mubin, "Predicting hospital stay length using explainable machine learning," IEEE Access, vol. 12, pp. 90571-90585, 2024.

[3]. M. A. Vollmer, B. Glampson, T. Mellan, S. Mishra, L. Mercuri, C. Costello, J. Cooke, S. L. Hegarty, A. Ayoubkhani, O. J. McInnes, K. M. Grint, H. J. T. Ward, A. Coupland, K. Bhaskaran, R. Williamson, E. J. Williamson, H. I. McDonald, L. Bates, N. Bacon, S. Walker, C. E. Tomlinson, K. Wing, A. J. Schultze, R. M. Andrews, C. Mehrkar, E. Herrett, B. Goldacre, L. Smeeth, B. MacKenna, and S. Bhatt, "A unified machine learning approach to time series forecasting applied to demand at emergency departments," BMC Emergency Medicine, vol. 21, no. 1, p. 9, 2021.

[4]. R. Gao, W. X. Cheng, P. N. Suganthan, and K. F. Yuen, "Inpatient discharges forecasting for Singapore hospitals by machine learning," IEEE Journal of Biomedical and Health Informatics, vol. 26, no. 10, pp. 4966-4975, 2022.

[5]. M. A. Bedada, I. Megdiche, E. Lamine, and D. Duval, "Multi-task learning approach for hospital bed requirement prediction," in 2024 IEEE/ACS 21st International Conference on Computer Systems and Applications (AICCSA), 2024, pp. 1-6.

[6]. 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.

[7]. Hamzaoui, A. Bouzir, and S. Benammou, "A comparative machine learning approaches for patient flow forecasting in an emergency department during the COVID-19," in 2022 14th International Colloquium of Logistics and Supply Chain Management (LOGISTIQUA), 2022, pp. 1-6.

[8]. T. I. Alshwaheen, Y. W. Hau, N. Ass'Ad, and M. M. Abualsamen, "A novel and reliable framework of patient deterioration prediction in intensive care unit based on long short-term memory-recurrent neural network," IEEE Access, vol. 9, pp. 3894-3918, 2020.

[9]. 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), 2024, pp. 1-7.

[10]. Yu, Y., Yao, S., Wang, K., Chen, Y., Wang, F. Y., & Lin, Y. (2023, October). Pursuing equilibrium of medical resources via data empowerment in parallel healthcare system. In 2023 IEEE International Conference on Systems, Man, and Cybernetics (SMC) (pp. 3186-3191). IEEE

[11]. J. W. Lin, W. Z. Wang, B. M. Zheng, C. H. Shih, H. M. Cheng, J. Hsu, and Y. C. Chen, "Artificial intelligence-driven hospital bed allocation for efficient patient admission," in International Conference on Applied System Innovation (ICASI 2025), vol. 2025, 2025, pp. 231-233.

[12]. P. Karthikeyan, M. Bharath, M. K. Vishnuram, and B. Sriram, "Forecasting patient length of stay for optimal hospital resource allocation," in 2024 IEEE International Conference on Smart Power Control and Renewable Energy (ICSPCRE), 2024, pp. 1-6.

[13]. R. Kashvi, V. Nirmala, and V. Vidhya, "Automated inter-hospital communication for bed allocation," in 2025 4th International Conference on Sentiment Analysis and Deep Learning (ICSADL), 2025, pp. 430-435.

[14]. T. H. Jamal, A. Y. Owda, and M. Owda, "Improve the efficiency of resource allocation in healthcare management using decision support system," in 2025 International Mobile, Intelligent, and Ubiquitous Computing Conference (MIUCC), 2025, pp. 467-472.

I. Hamzaoui, A. Bouzir, and S. Benammou, "A dual LSTM-based framework for forecasting patient flow and classifying COVID-19 severity in emergency departments: A case study from

Tunisia," in 2025 16th International Conference on Logistics and Supply Chain Management (LOGISTIQUA), 2025, pp. 1-6.