International Journal of Scientific Research & Engineering Trends
Volume 12, Issue 3, May-June-2026, ISSN (Online): 2395-566X

MediCast: Smart Hospital ICU Beds and Oxygen Demand Predictor

Sahil Arun Sahane, Suhani Sharma, Amay Prasad Sabnis,
Suhani Singh, Assistant Professor Rahul B. Mandlik
Computer Engineering
MET BKC Institute Of Engineering, Nashik, India

Abstract- Efficient management of critical hospital resources such as intensive care unit (ICU) beds, oxygen supply, and medical staff has become a major challenge, particularly during large-scale healthcare emergencies. Conventional hospital management systems are largely reactive and often fail to anticipate sudden surges in patient demand, resulting in delayed responses and resource shortages. This paper presents MediCast, an AI-driven hospital resource forecasting and decision support system designed to predict ICU bed occupancy and oxygen demand in advance while supporting optimized resource allocation. The proposed framework employs Long Short-Term Memory (LSTM) networks for time-series forecasting of ICU admissions and oxygen consumption trends, and XGBoost models for learning complex patterns from structured hospital data. Based on the predicted demand, an optimization layer assists in efficient allocation of beds and staff resources to reduce overload and improve preparedness. The system also provides an interactive dashboard for real-time visualization of predictions, alerts, and analytical insights, enabling hospital administrators to take proactive decisions. By integrating predictive analytics and optimization within a unified platform, MediCast enhances operational efficiency, minimizes critical resource shortages, and supports data-driven healthcare management in high-demand scenarios.

Keywords— ICU bed prediction, oxygen demand forecasting, artificial intelligence in healthcare, LSTM, XGBoost, hospital resource management, predictive analytics

I. INTRODUCTION

Efficient management of ICU beds and oxygen supply is a critical challenge for hospitals, especially during sudden increases in patient admissions. In emergency situations, limited infrastructure and delayed decision-making often lead to resource shortages and operational stress. Most traditional hospital management systems focus on data recording and administrative tasks, but they lack predictive capabilities to anticipate future demand.

Recent advancements in artificial intelligence provide an opportunity to improve healthcare planning through data driven forecasting. Machine learning models can analyze historical hospital data to identify patterns in patient inflow, ICU occupancy, and oxygen consumption. By using time-series and structured data analysis, hospitals can move from reactive responses to proactive preparation.

To address this need, this paper proposes MediCast, an AI based hospital resource forecasting framework. The system combines XGBoost and LSTM models to predict ICU bed occupancy and oxygen demand, along with a Mixed-Integer Linear Programming optimization model for efficient resource allocation. A web-based dashboard further supports real-time monitoring and decision-making. The proposed approach aims to enhance hospital preparedness and improve operational efficiency during high-demand scenarios.

II. LITERATURE REVIEW

The rapid adoption of artificial intelligence (AI) and machine learning (ML) techniques in healthcare has significantly improved predictive analytics, clinical decision support, and hospital resource management. Various studies have explored AI-driven approaches for forecasting patient admissions, optimizing staff allocation, and managing critical resources such as ICU beds and oxygen supply.

Several works have focused on workforce optimization using AI-based models. Zhang et al. [1] proposed an AI-guided stochastic staffing framework to improve nurse scheduling efficiency and reduce operational costs. Similarly, Yinusa and Faezipour [4] utilized Mixed-Integer Linear Programming (MILP) for optimizing hospital resource allocation, demonstrating improved workflow efficiency. However, these approaches primarily focus on staffing and lack integration with predictive demand forecasting.

Oxygen demand prediction has been widely studied, especially during the COVID-19 pandemic. Tempe et al. [2] developed mathematical models to estimate oxygen consumption based on patient load and treatment duration. The World Health Organization [7] also provided guidelines and reports on oxygen distribution systems during healthcare emergencies. While these studies provide useful baseline insights, they lack real-time adaptability and predictive intelligence.

Time-series and deep learning techniques have shown promising results in healthcare forecasting. Choi et al. [3] introduced an LSTM-based model for ICU admission prediction, achieving high accuracy in identifying sequential patterns. Similarly, Huser et al. [6] demonstrated the effectiveness of LSTM models for early detection of respiratory failure in ICU patients. Wang et al. [13] further extended deep learning approaches for predicting patient deterioration using multivariate time-series data. Although these models are effective, they are primarily focused on patient-level prediction rather than hospital-wide resource planning.

In addition to forecasting, optimization techniques such as genetic algorithms and reinforcement learning have been applied to healthcare resource management. Karboub and Tabaa [5] proposed a genetic algorithm-based approach for patient bed allocation, improving utilization efficiency. Song et al. [12] introduced a reinforcement learning framework for hospital bed capacity management. However, these approaches lack integration with predictive models and are not designed for dynamic real-time scenarios.

Recent studies also highlight the importance of integrating machine learning with healthcare systems. Rajkomar et al. [10] discussed the role of machine learning in modern medicine, emphasizing its potential in improving clinical decision-making. Miotto et al. [8] provided a comprehensive review of deep learning applications in healthcare, outlining opportunities and challenges. Chen et al. [9] explored predictive analytics for hospital resource management, while Rahman et al. [11] demonstrated machine learning-based prediction of ICU admissions and mortality. Alvarez et al. [14] further proposed hybrid frameworks combining optimization and machine learning for healthcare planning.

Despite significant advancements, existing approaches remain fragmented, addressing prediction, optimization, or visualization independently. There is a lack of a unified framework that integrates forecasting, optimization, and real-time decision support. The proposed MediCast system addresses this gap by combining XGBoost and LSTM-based forecasting with MILP driven optimization into a single integrated platform, enabling proactive and data-driven hospital resource management.

III. PROPOSED SYSTEM

The proposed system, MediCast: Smart Hospital ICU Beds and Oxygen Predictor, is designed to address the critical shortcomings of conventional hospital management systems that rely heavily on manual planning, static rules, and isolated departmental workflows. During emergency situations such as pandemics or seasonal disease outbreaks, these traditional systems fail to anticipate sudden surges in ICU admissions, oxygen consumption, and staffing requirements, resulting in delayed decision-making and inefficient resource utilization. MediCast introduces an integrated, AI-driven framework that enables proactive hospital resource planning through accurate forecasting and intelligent optimization.

MediCast follows a layered architecture that combines predictive analytics, optimization models, and real-time visualization into a unified decision-support system. The framework ingests historical and real-time hospital data, including patient admission records, ICU occupancy logs, oxygen usage statistics, and staff schedules. This data is preprocessed to handle missing values, normalize temporal variations, and extract relevant features required for model training and inference. The processed data then flows into the predictive modeling layer, which consists of machine learning and deep learning models tailored for different forecasting tasks.

For ICU admission prediction, MediCast employs the XGBoost algorithm due to its robustness in handling structured healthcare data and its ability to capture nonlinear relationships between patient inflow patterns, demographic variables, and historical admission trends. XGBoost generates short- to medium-term forecasts of ICU bed demand, enabling hospitals to anticipate occupancy saturation points. In parallel, Long Short-Term Memory (LSTM) networks are utilized to model temporal dependencies in oxygen consumption data. LSTM's gated architecture allows it to effectively learn long term trends and seasonal patterns in oxygen usage, making it suitable for forecasting continuous demand under varying clinical conditions.

While predictive models estimate future demand, MediCast also incorporates a prescriptive analytics layer to optimize hospital operations. Mixed-Integer Linear Programming (MILP) is used to formulate resource allocation and staff scheduling problems as constrained optimization tasks. The

MILP model considers multiple constraints, including ICU bed capacity, oxygen availability, staff shift limits, patient-to-staff ratios, and priority-based resource allocation.

A key strength of MediCast lies in the seamless integration between forecasting and optimization components. Predicted ICU admissions and oxygen requirements serve as direct inputs to the MILP model, ensuring that optimization decisions are informed by future demand rather than reactive assumptions. This closed-loop design allows the system to continuously adapt as new data becomes available, improving

Table1: Comparative analysis

Study	Approach	Limitations	Findings
Si Zhang et al.	Staffing optimization	No demand prediction	Reduced cost
Deepak Tempe et al.	Oxygen modeling	Static assumptions	Baseline model
Sooho Choi et al.	LSTM prediction	Limited scope	High accuracy
Yinusa et al.	MILP optimization	Cost-focused	Efficient allocation
Karboub et al.	Genetic algorithm	No forecasting	Better utilization
Huser et al.	LSTM detection	Small dataset	Early detection

both prediction accuracy and operational efficiency over time. To ensure practical usability, MediCast includes a web based dashboard developed using frameworks such as Streamlit or Flask. The dashboard presents forecast trends, optimization outputs, and system alerts through intuitive visualizations, enabling hospital administrators to monitor key performance indicators in real time. When the system detects an impending resource shortage—such as ICU saturation or oxygen depletion—it triggers early alerts, allowing administrators to initiate corrective actions well in advance.

Overall, MediCast represents a shift from reactive hospital management to proactive, data-driven decision-making. By integrating machine learning–based forecasting, deep learning–driven temporal analysis, and mathematical optimization within a single framework, the system enhances hospital preparedness, reduces operational risk, and supports timely, informed decisions during critical healthcare scenarios.

Objectives

The main objectives of the proposed system are as follows:

- To study the existing hospital resource management systems and identify their limitations, focusing on the unavailability of predictive and optimization mechanisms during crises.
- To forecast ICU bed requirements and oxygen demand using machine learning algorithms like XGBoost and LSTM.
- To optimize the scheduling and allocation of limited resources such as staff, ICU beds, and oxygen cylinders using MILP-based optimization.
- To develop a real-time alert and visualization dashboard that assists administrators in proactive planning and decision-making.
- To deploy the system as a user-friendly web application that integrates prediction, optimization, and alerting into a single unified interface.

System Workflow

The workflow of the proposed MediCast system follows a structured and sequential process to ensure accurate prediction and efficient management of resources:

Data Collection: Data is collected from verified hospital information systems, open-source repositories, and simulated datasets. The collected data includes hospital admission records, ICU bed occupancy details, oxygen usage logs, and staff shift schedules.

Data Preprocessing: The data is cleaned to remove inconsistencies, handle missing values, and normalize numerical fields. Categorical features such as patient type, diagnosis code, or oxygen unit category are encoded to ensure compatibility with machine learning models.

Prediction using Machine Learning Models: The preprocessed data is divided into training and testing subsets. Two complementary machine learning models are applied: XGBoost, which predicts ICU occupancy based on infection rate, patient admission trends, and hospital inflow; and LSTM (Long Short-Term Memory), which analyzes time-series oxygen consumption data to capture sequential dependencies and daily usage fluctuations. Together, they generate accurate short-term and long-term forecasts for ICU demand and oxygen supply.

Resource Optimization: The optimization module, implemented using Mixed-Integer Linear Programming (MILP), utilizes the predicted data to compute efficient resource allocation plans. For instance, when predicted oxygen usage exceeds capacity, MILP suggests optimized redistribution strategies prioritizing critical patients. It also assists in staff scheduling to balance workload and improve operational efficiency.

Dashboard Integration: The outputs from the prediction and optimization modules are integrated into a realtime dashboard developed using Streamlit or Flask. The dashboard visualizes ICU occupancy forecasts, oxygen stock levels, and staff schedules through interactive charts and tables.

Alerts and Decision Support: The system provides automated alerts and notifications when resource thresholds are predicted to be breached. This enables administrators to take proactive measures and ensures end-toend automation and data-driven decision support.

Dataset Description

The MediCast system is evaluated using a combination of publicly available healthcare datasets and simulated hospital operational data to reflect real-world ICU and oxygen usage scenarios. The dataset includes daily hospital admission records, ICU occupancy counts, oxygen consumption logs, and staff scheduling information collected over multiple months.

The primary features include patient admission rate, ICU length of stay, oxygen flow rate per patient, recovery duration, and staff-to-patient ratios. Time-stamped records enable sequential modeling of oxygen usage trends. Since real-time hospital data access is restricted due to privacy constraints, simulated datasets were generated following realistic statistical distributions derived from published healthcare reports and prior studies.

All datasets were anonymized and aggregated at the hospital level, ensuring no personally identifiable patient information was used during model training or evaluation.

System Architecture and Component Details

The proposed system architecture follows a modular, multilayered pipeline designed for scalability, automation, and realtime hospital resource forecasting. Each layer performs a distinct function — from data acquisition to predictive modeling, optimization, and visualization — ensuring transparency, accuracy, and operational efficiency.

- Data Collection Layer. This layer handles the acquisition of hospital-related data from multiple verified sources. It integrates:
- Data Preprocessing Layer. The preprocessing layer cleans, filters, and standardizes hospital data before analysis. It performs:
- Removal of missing, duplicate, or inconsistent entries.
- Label and one-hot encoding for categorical variables such as department, diagnosis, or patient type.
- Normalization of numerical values including oxygen consumption rate, admission count, and ICU occupancy ratio.
- Feature engineering to derive new parameters like average patient stay duration and daily oxygen usage per patient.

This step ensures the dataset is structured and reliable for model training and optimization.

Forecasting Layer. This module predicts ICU bed demand and oxygen consumption using hybrid AI models:

- XGBoost: Used for predicting ICU occupancy and patient inflow based on non-sequential features such as infection rate and admission trends.
- LSTM: A deep learning time-series model designed to capture temporal patterns and daily fluctuations in oxygen usage.
- Forecasting Horizon: Produces short-term predictions (typically 7 days) with uncertainty intervals to account for data variability.

The forecasting results serve as inputs for the next stage—ICU and oxygen modeling.

ICU Beds and Oxygen Modeling Layer. This layer translates predictive outputs into meaningful operational metrics:

- Calculates estimated ICU occupancy levels and patient turnover rates.
- Models oxygen consumption patterns based on forecasted patient load.
- Provides actionable data to the optimization layer for planning and allocation.

It acts as a bridge between predictive analytics and real-world hospital operations.

Optimization and Scheduling Layer. This layer uses Mixed-Integer Linear Programming (MILP) to optimize hospital resource allocation. It considers:

- Constraints such as bed availability, oxygen stock, staff skills, and shift duration.
- Objectives like minimizing oxygen wastage and balancing staff workload.
- Real-time recommendations such as redistributing oxygen cylinders or rescheduling medical personnel.

The optimization process ensures that hospitals maintain maximum efficiency under limited resources.

Alerts and Dashboard Layer. This is the system's interactive visualization and communication component. It provides:

- A web-based dashboard (built using Streamlit or Flask) for real-time monitoring of forecasts and optimization outcomes.
- Visual analytics such as ICU bed occupancy graphs, oxygen consumption curves, and staff workload charts.
- Automated alert notifications via email or messaging when ICU occupancy exceeds 95% or oxygen supply falls below two days.

The dashboard enables hospital administrators to make fast, data-driven decisions.

Evaluation and Deployment Layer. This layer tests, validates, and deploys the trained models. It includes:

- Model evaluation using MAE, RMSE, and accuracy metrics.
- Continuous integration with hospital data for live predictions.
- Secure deployment using Docker or Flask API to enable real-time access.

The layer ensures that the MediCast system remains robust and production-ready.

Continuous Improvement and Retraining Layer. MediCast includes an adaptive feedback mechanism for selflearning and model refinement:

- Compares predicted versus actual ICU occupancy and oxygen usage to evaluate model drift.
- Periodically retrains forecasting models with new hospital data.
- Incorporates administrator feedback to improve alert accuracy and usability.

This makes MediCast a dynamic system capable of evolving with changing healthcare conditions.

Overall, the architecture ensures a smooth and transparent flow from raw hospital data to intelligent predictions, optimization, and decision support. Its modular structure guarantees scalability, reliability, and continuous improvement for effective hospital resource management.

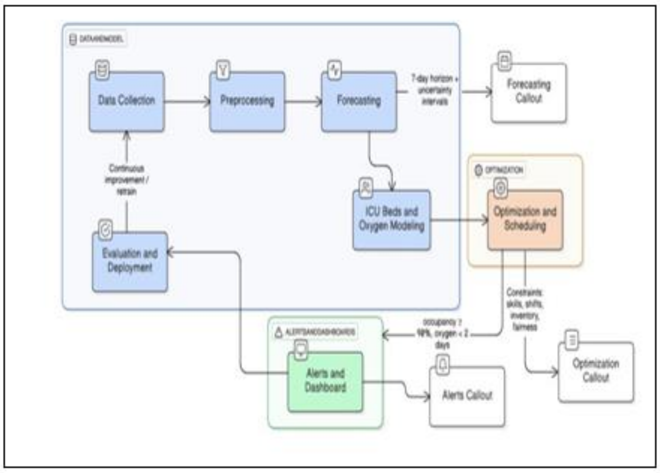

Mathematical Formulation of the Optimization Model

Hospital resource allocation in MediCast is formulated as a Mixed-Integer Linear Programming (MILP) problem. The objective is to optimally allocate ICU beds, oxygen supply, and medical staff while minimizing shortages and workload imbalance. The model operates over discrete time intervals using predicted demand from forecasting models. Decision Variables:

- xb(t): ICU beds allocated at time t
- xo(t): Oxygen allocated at time t
- xs(t): Staff assigned at time t
- ub(t): ICU bed shortage at time t
- uo(t): Oxygen shortage at time t Objective Function:
- $\min X(\alpha uo(t) + \beta ub(t) + \gamma|xs(t) - \bar{x}s|)$ (1)

t

Constraints:

Capacity Constraints:

$xb(t) \leq Bmax$, $xo(t) \leq Omax$, $xs(t) \leq Smax$ (2)

Demand Constraints:

$xb(t) + ub(t) \geq db(t)$ (3) $xo(t) + uo(t) \geq do(t)$ (4)

Non-negativity:

$xb(t),xs(t) \in Z+$, $xo(t),ub(t),uo(t) \geq 0$ (5)

This formulation ensures optimal and feasible allocation of hospital resources by balancing predicted demand with system constraints.

IV. DISCUSSION

The proposed MediCast system shows how predictive analytics and optimization can improve hospital resource management. Unlike traditional methods that rely on manual estimation, this framework forecasts ICU bed demand and oxygen usage in advance, allowing administrators to prepare before shortages occur.

By combining XGBoost for ICU prediction and LSTM for oxygen forecasting, the system captures both structured patterns and time-based trends. The MILP optimization module then converts these predictions into practical allocation strategies while considering real-world constraints such as bed capacity and staff limits.

Although the evaluation is based on simulated and public datasets, the results indicate that the approach can support proactive planning during high-demand situations. Overall, MediCast demonstrates the potential of integrating AI and optimization to enhance hospital preparedness and operational efficiency.

V. CONCLUSION

This paper presented MediCast, an intelligent and integrated framework for hospital resource forecasting and operational planning. By combining machine learning–based prediction with mathematical optimization, MediCast addresses critical challenges faced by healthcare institutions during high-demand scenarios. The system leverages XGBoost for forecasting ICU bed occupancy and Long Short-Term Memory (LSTM) networks for modeling temporal oxygen demand, while a Mixed-Integer Linear Programming (MILP) formulation ensures optimal allocation of ICU beds, oxygen resources, and medical staff under operational constraints.

The proposed approach enables hospitals to transition from reactive crisis management to proactive, data-driven preparedness. The integration of predictive analytics with optimization allows MediCast to anticipate resource shortages, balance staff workloads, and improve overall utilization efficiency. Furthermore, the inclusion of a real-time, web-based dashboard and automated alert mechanisms enhances situational awareness and supports timely decision-making by hospital administrators.

MediCast is designed to be modular, scalable, and adaptable, making it suitable for deployment across healthcare facilities of varying sizes and capacities. The framework can be easily extended to incorporate additional predictive components such as ventilator demand forecasting, patient discharge estimation, or inter-hospital resource coordination. Future work will focus on integrating real-time electronic health record (EHR) streams, validating the system using large-scale real-world hospital datasets, and exploring advanced deep learning and federated learning techniques to support regional and national-level healthcare planning.

Overall, MediCast demonstrates the potential of artificial intelligence and optimization-driven decision support to enhance hospital resilience, improve resource efficiency, and ultimately contribute to better patient care outcomes in modern healthcare systems.

REFERENCES

1. S. Zhang et al., “AI-guided nurse staffing and scheduling in healthcare systems,” IEEE Access, vol. 13, pp. 22045–22058, 2025.
2. D. Tempe et al., “Mathematical modeling of hospital oxygen consumption during COVID-19,” J. Med. Syst., vol. 49, no. 2, pp. 180–189, 2024.
3. S. Choi et al., “Sequential anomaly detection for ICU admission prediction using LSTM,” Comput. Biol. Med., vol. 165, Art. no. 107421, 2024.
4. A. Yinusa and M. Faezipour, “Optimization of hospital staffing and resource allocation using MILP,” Healthcare Technol. Lett., vol. 10, no. 1, pp. 15–25, 2023.
5. K. Karboub and M. Tabaa, “Genetic algorithm approach for patient bed allocation optimization,” Procedia Comput. Sci., vol. 225, pp. 301–309, 2023.
6. M. Huser¨ et al., “LSTM-based early respiratory failure detection in ICU patients,” Artif. Intell. Med., vol. 115, Art. no. 102142, 2021.
7. World Health Organization, Oxygen Sources and Distribution for COVID-19 Treatment, WHO Technical Report, Geneva, Switzerland, 2021.
8. R. Miotto et al., “Deep learning for healthcare: Review, opportunities and challenges,” Brief. Bioinform., vol. 19, no. 6, pp. 1236–1246, 2018.
9. J. Chen et al., “Predictive analytics for hospital resource management,” IEEE J. Biomed. Health Inform., vol. 26, no. 4, pp. 1812–1821, 2022.
10. A. Rajkomar et al., “Machine learning in medicine,” N. Engl. J. Med., vol. 380, no. 14, pp. 1347–1358, 2019.
11. M. Rahman et al., “Predictive modeling of ICU admission and mortality using machine learning,” Comput. Methods Programs Biomed., vol. 196, Art. no. 105582, 2021.
12. X. Song et al., “A reinforcement learning approach for hospital bed capacity management,” IEEE Access, vol. 10, pp. 112345–112356, 2022.
13. Y. Wang et al., “Early prediction of patient deterioration using deep learning on multivariate time series,” IEEE J. Biomed. Health Inform., vol. 24, no. 1, pp. 260–268, 2020.

14. J. Alvarez et al., “Hybrid optimization and machine learning framework for hospital resource planning,” Expert Syst. Appl., vol. 213, Art. no. 119107, 2023.