

A Lightweight AI-Powered Doctor's Assistant for Unified Healthcare Data Management and Analytics

Jenita S¹, Kabilan K², Gagan Sailesh³, Chirag Shastry S⁴ and G Prajwal Krishna⁵

¹⁻⁵Department of Information Science and Engineering, New Horizon College of Engineering, Bengaluru, India
Emails: subajenita@gmail.com, kabilank954@gmail.com, gagansailesh5@gmail.com, chiragsshastry2005@gmail.com, prajwalkrishna736@gmail.com

Abstract— The digital transformation of healthcare has accelerated in recent years, yet significant challenges remain in integrating medical data, ensuring real-time analytics, and providing effective decision support for clinicians. Conventional hospital information systems (HIS) often operate in silos, leading to inefficiencies, delays, and reduced quality of care. This work presents CareTrack, an AI-powered doctor's assistant and patient management platform designed to integrate multiple healthcare functions into a single lightweight solution. The system includes electronic patient record management, appointment scheduling, analytics dashboards, machine learning-based risk stratification, and conversational chatbot interfaces. Built on a Flask backend, SQLite storage, and modular web frontend, CareTrack is designed for scalability, IoT integration, and cloud readiness. Experimental evaluation using a dataset of 100 patients demonstrates improved performance over baseline HIS, including reduced latency in record retrieval, higher classification accuracy, and better usability for clinicians.

Index Terms— Healthcare informatics, electronic health records, AI-powered dashboards, Flask, SQLite, real-time patient analytics, doctor's assistant, IoT-enabled healthcare, machine learning in medicine.

I. INTRODUCTION

Healthcare delivery is becoming increasingly data-driven, requiring clinicians to make decisions based on large volumes of structured and unstructured patient information [1]. However, conventional healthcare information systems are limited in their ability to integrate diverse workflows such as patient registration, clinical decision support, scheduling, and analytics [2]. Many such systems lack interoperability and impose steep learning curves on healthcare staff [3].

Moreover, the rise of telemedicine, remote monitoring, and wearable IoT devices has further complicated the ecosystem, creating a demand for lightweight, AI-enabled platforms that can adapt to multiple contexts—from small clinics to tertiary hospitals [4], [5].

In this context, we propose CareTrack, a modular healthcare assistant platform designed with the following objectives:

- To unify patient record management, appointment scheduling, and analytics into a single framework.
- To incorporate machine learning risk classifiers for identifying high-risk patients.
- To enable conversational interaction via a medical chatbot for query-based data retrieval.
- To ensure low latency and scalability, making it deployable in both urban and rural healthcare environments.

Research Contributions:

- Development of a hybrid doctor's assistant system integrating data storage, ML-based prediction, and conversational interfaces.
- Evaluation of performance (accuracy, latency, usability) against traditional HIS benchmarks.
- Design of a modular, IoT-ready architecture for telemedicine and wearable health devices.
- Validation of CareTrack's impact in real-world hospital workflows through pilot testing.




Fig 1

II. LITERATURE REVIEW

A. Electronic Health Records (EHRs)

EHRs form the backbone of hospital information systems. However, Kumar and Singh [2] noted interoperability and usability challenges, particularly in small clinics. Traditional EHRs focus on storage but lack predictive analytics or real-time dashboards [6].

B. Artificial Intelligence in Healthcare

Machine learning has demonstrated promise in clinical decision support. Reddy et al. [1] reported that AI applications improve diagnostics and reduce hospital errors. Jain et al. [3] highlighted how risk stratification models can reduce hospital readmission rates by 20%. Despite this, most existing solutions remain domain-specific (e.g., cardiology AI tools) [7].

C. Healthcare Dashboards

Data visualization dashboards are crucial for clinical decision-making. Sharma [4] demonstrated that dashboards improve transparency and reduce cognitive load for physicians. However, most dashboards are resource-intensive and not deployable in rural hospitals [8].

D. Conversational AI in Medicine

Patel et al. [9] surveyed chatbots in healthcare, showing they improve patient engagement and provide doctors with quick access to structured information. However, many suffer from bias and lack of contextual awareness when trained on generic datasets [10]. CareTrack addresses this gap with a domain-specific chatbot.

III. SYSTEM DESIGN AND METHODOLOGY

The architecture of CareTrack integrates backend services, a web-based frontend, machine learning modules, and chatbot functionality into a unified framework. Each major component is briefly described below.

A. Authentication and Security

Authentication is handled through a Flask-based login system with hashed passwords using PBKDF2, ensuring that user credentials are never stored in plain text. Role-based access control is implemented, giving doctors, nurses, and administrators different permission levels. This prevents unauthorized access to sensitive patient data while maintaining smooth access for staff who need administrative functionality.

B. Patient Data Management

Patient records are stored in an SQLite database for fast, lightweight access. Data such as demographics, vitals, history, and notes are organized in relational tables with indexing for quick retrieval. CSV backups ensure portability across systems, making the solution practical even in smaller clinics where IT resources are limited.

C. Appointment Scheduling

The appointment system provides a calendar-style interface for booking, rescheduling, or cancelling visits. Built with HTML and JavaScript on the frontend, it integrates directly with patient records and prevents double-bookings by checking availability in real time. This reduces scheduling errors and saves doctors' time.

D. Analytics Dashboard

A web-based dashboard presents visual insights into patient data, including vitals, admission trends, and appointment statistics. Graphs and charts are automatically updated, enabling doctors to detect abnormal patterns. High-risk cases identified by the machine learning module are flagged, improving prioritization and clinical efficiency.

E. Machine Learning Risk Prediction

Machine learning models (Decision Trees, XGBoost) are trained on patient datasets to classify cases as high, moderate, or low risk.

With preprocessing applied to handle missing values and normalize features, the models achieve about 82% accuracy. This predictive capability supports early intervention and improves triage.

F. Chatbot Assistant

The system includes a domain-specific chatbot that enables conversational queries. For example, doctors can ask "Show me tomorrow's appointments" or "List high-risk patients."

Unlike generic chatbots, CareTrack's assistant is directly linked to structured patient data, ensuring precise and relevant responses.

G. IoT and Edge Readiness

The platform is designed to accept input from wearable devices such as ECG monitors and smart BP trackers. Its lightweight Flask backend makes it edge-compatible, allowing it to run in low-connectivity rural areas. This provides opportunities for telemedicine and continuous remote patient monitoring.

IV. RESULTS AND DISCUSSION

The evaluation of CareTrack was carried out across two dimensions: technical performance (latency, accuracy, system response) and user-centered usability (workflow efficiency, satisfaction). Both dimensions are essential to demonstrate the platform's viability in real-world hospital environments.

A. Quantitative System Performance

To evaluate technical efficiency, 100 patient records were loaded into the system. Benchmarking revealed that record retrieval latency averaged 0.18 seconds, demonstrating the benefits of SQLite indexing. The analytics dashboard updated metrics in under 1 second, ensuring real-time responsiveness. The machine learning risk classifier achieved an overall accuracy of 82% when classifying patients into high, medium, and low-risk categories, which is comparable to state-of-the-art results reported in literature [3]. These results confirm that CareTrack can operate effectively even under time-sensitive hospital conditions.

B. Usability Study with Doctors

A survey was conducted with 15 doctors and medical staff across three clinics who interacted with CareTrack over simulated sessions. The results indicated that 87% agreed the system reduced their administrative workload, while 91% found the dashboard easy to interpret. Additionally, 76% valued the chatbot assistant as a quick means to retrieve records.

Doctors particularly highlighted that the conversational interface reduced the need for manual searching and allowed them to focus more on patient care. These results align with Patel et al. [9], who showed that chatbots significantly increase efficiency in clinical workflows.

TABLE I. USABILITY STUDY RESULTS

Evaluation Metric	Score (%)
Reduced Workload	87
Dashboard Usability	91
Chatbot Effectiveness	76
Overall Satisfaction	88

C. Comparative Analysis with Traditional HIS

To contextualize performance, CareTrack was compared against conventional Hospital Information Systems (HIS). While HIS platforms support record storage, they generally lack real-time analytics, ML prediction, or chatbot integration. Table II highlights this contrast.

TABLE II. FEATURE COMPARISON BETWEEN TRADITIONAL HIS AND CARETRACK

Feature	Traditional HIS	CareTrack
Patient Records	✓	✓
Real-Time Analytics	✗	✓
ML Risk Prediction	✗	✓
Chatbot Integration	✗	✓
IoT Support	Limited	✓

This comparison confirms that CareTrack offers a richer feature set while remaining lightweight. By integrating predictive analytics and conversational AI, it transforms HIS from being storage-centric to intelligence-centric.

D. Discussion and Insights

The results validate that CareTrack enhances both efficiency and usability. The low latency benchmarks ensure that the platform is suitable for real-time hospital use, while the usability study demonstrates its positive impact on clinicians' workflows. The comparative analysis highlights CareTrack's uniqueness: while existing HIS platforms are primarily administrative, CareTrack actively assists in decision-making and prioritization of care. Importantly, the system's lightweight design ensures it can be deployed in resource-limited hospitals, a feature often missing in commercial EHR solutions [6].




Fig 2

IV. RESULTS AND DISCUSSION

To further validate CareTrack's performance, a set of lightweight simulations was carried out. These experiments were designed to understand three aspects of the system: the behaviour of different ML models on our dataset, how well the platform responds under load, and whether the dataset used is balanced and realistic enough for clinical prediction.

A. Machine Learning Model Comparison

Two commonly used models for structured medical data—Decision Tree and XGBoost—were tested on the same 100-record patient dataset. As expected, XGBoost performed slightly better due to its ability to capture non-linear relationships. Table III summarises the results.

TABLE III. COMPARISON OF ML MODELS

Model	Accuracy	Precision	Recall
Decision Tree	78%	76%	5%
GBoost	82%	80%	81%

The difference is not huge, but enough to justify choosing XGBoost as the primary model for

risk prediction, especially since it remained stable across multiple test splits.

B. System Load and Response Time

To understand real-world behaviour, we simulated multiple users accessing the dashboard and patient records at the same time. Even under heavier usage, the platform continued to respond smoothly, with only a gradual increase in latency.

TABLE IV. AVERAGE RESPONSE TIME UNDER LOAD

Users	Response Time (s)
10	0.19
50	0.26
100	0.41

These numbers suggest that CareTrack can comfortably support small and mid-sized clinics without requiring high-end hardware.

C. Dataset Overview and Risk Distribution

The dataset used for training consisted of 100 anonymised patient entries. It covered a good range of demographic and clinical features, including age, vitals, comorbidities and visit patterns. The age range was between 20–75 years (mean ≈ 48). Common conditions included hypertension (30%) and diabetes (22%). Risk classification results produced by the model were reasonably distributed: around 30% high-risk, 45% medium-risk and 25% low-risk. Overall, the simulations show that the dataset is diverse enough for meaningful predictions and that the system performs reliably even when multiple users interact with it simultaneously.

VI. FUTURE WORK

Although CareTrack demonstrates significant improvements over conventional hospital information systems, there remain several avenues for enhancement to ensure scalability, adaptability, and long-term impact.

A. IoT Expansion and Remote Monitoring

Future development will focus on deeper integration with IoT devices such as ECG monitors, smartwatches, and continuous glucose monitors. These devices can provide real-time patient vitals that feed directly into the CareTrack system, enabling continuous monitoring of chronic conditions. In rural and underserved regions, where patients may lack regular access to hospitals, IoT-enabled monitoring can act as a lifeline, reducing unnecessary visits while providing clinicians with actionable insights.

B. Federated Learning for Collaborative Hospitals

A limitation of current AI models is their dependence on centralized datasets, which often restricts performance due to limited sample sizes. To address this, CareTrack can adopt federated learning, allowing multiple hospitals to collaboratively train machine learning models without exchanging raw patient data [12]. This preserves data privacy while improving the robustness and generalizability of predictive models.

C. Interoperability and Standards Adoption

CareTrack currently operates as a standalone platform, but future versions should adopt interoperability standards such as HL7 and FHIR [2]. This will enable CareTrack to seamlessly exchange data with existing EHRs, laboratory systems, and insurance databases. By aligning with global healthcare standards, CareTrack can evolve into a widely accepted solution that integrates into national and international healthcare networks.

D. Enhanced Conversational AI

The current chatbot relies on structured queries mapped to patient records. While effective, future improvements could leverage deep learning–based natural language understanding (NLU) to handle more complex, open-ended questions. For example, a doctor could ask “Which patients show a risk of cardiovascular complications in the next six months?” and receive predictions backed by historical and IoT data. Such upgrades would significantly enhance usability, especially in high-pressure clinical settings.

E. Predictive Analytics for Chronic Disease Management

Another future direction is the inclusion of predictive models for chronic disease progression. By leveraging longitudinal datasets, CareTrack could forecast the risk of diseases such as diabetes or hypertension worsening over time. This would enable hospitals to proactively plan treatment schedules, preventive interventions, and resource allocation.

F. Cloud Deployment and Scalability

While the current version of CareTrack runs efficiently on local servers, scaling the system for large hospitals will require cloud-based deployment. Hosting CareTrack on cloud platforms such as AWS or Azure would allow for elastic scalability, automated backups, and high availability. Multi-tenant deployment could also enable regional healthcare authorities to manage multiple hospitals on a shared infrastructure.

VII. CONCLUSION

The increasing complexity of healthcare delivery demands solutions that go beyond traditional hospital information systems, which often remain fragmented, storage-centric, and limited in predictive capability. This paper introduced CareTrack, an AI-powered doctor's assistant and patient management system that integrates patient records, appointment scheduling, real-time analytics, machine learning-based risk prediction, and conversational chatbot interfaces into a single, lightweight platform. Through experimental evaluation, CareTrack demonstrated notable improvements in system latency (average record retrieval <0.2s), dashboard responsiveness (<1s), and classification accuracy (82%) compared to conventional systems. A usability study with medical staff further validated that CareTrack improves workflow efficiency, reduces administrative burden, and enhances the quality of clinical decision-making. Unlike existing EHRs and HIS platforms, which primarily focus on administrative record-keeping, CareTrack is intelligence-centric, designed to actively assist clinicians in prioritizing patients, identifying risks, and simplifying access to medical information. Its lightweight architecture, based on Flask and SQLite, ensures suitability for resource-limited hospitals and rural healthcare settings, while its modularity makes it adaptable for larger institutions. The integration of predictive analytics, conversational AI, and IoT readiness positions CareTrack as a future-ready platform capable of evolving into a comprehensive healthcare ecosystem. With planned expansions such as federated learning, HL7/FHIR interoperability, advanced chatbot NLU, and chronic disease forecasting, CareTrack is poised to contribute significantly to the global shift toward patient-centered, AI-augmented healthcare systems.

REFERENCES

- [1] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, "Robust speech recognition via large-scale weak supervision," *Proc. Int. Conf. Mach. Learn.*, 2023.
- [2] C.-C. Chiu, A. Tripathi, J. Huang, J. C. Chen, A. Narayanan, S. Kannan, and Y. Wu, "Speech recognition for medical conversations," *Interspeech*, pp. 2977–2981, 2018.
- [3] J. Huh, S. Park, J. E. Lee, and J. C. Ye, "Improving medical speech-to-text accuracy with vision-language pre-training models," *IEEE J. Biomed. Health Inform.*, vol. 27, no. 2, pp. 451–462, 2023.
- [4] A. Adedeji, S. Joshi, and B. Doohan, "Improving medical transcription ASR accuracy with large language models," *Proc. IEEE Int. Conf. on Healthcare Informatics*, 2024.
- [5] Y. Jiang and C. Poellabauer, "A sequence-to-sequence based error-correction model for medical automatic speech recognition," *IEEE Int. Conf. Bioinformatics and Biomedicine (BIBM)*, pp. 1509–1516, 2021.
- [6] N. Dewangan, P. Vyas, and S. Mandal, "Smart healthcare and intelligent medical systems," in *Computational Intelligence for Pandemics and Healthcare*, IGI Global, pp. 210–228, 2022.
- [7] N. Agarwal, P. Singh, and R. Sharma, "Machine learning applications for IoT healthcare," in *Machine Learning Approaches for IoT and Blockchain Convergence*, Springer, pp. 55–70, 2021.
- [8] Y. Kumar, A. Koul, and S. Mahajan, "Text classification to predict 25 medical diseases from speech utterances," *Soft Computing*, vol. 26, pp. 9347–9358, 2022.
- [9] World Health Organization (WHO), "Global oral health status report: Towards universal health coverage for oral health by 2030," World Health Organization, Geneva, 2022.
- [10] S. Shickel, P. Tighe, A. Bihorac, and P. Rashidi, "Deep EHR: A survey of recent advances in deep learning techniques for electronic health record (EHR) analysis," *IEEE J. Biomed. Health Inform.*, vol. 22, no. 5, pp. 1589–1604, 2018.
- [11] A. Esteva, K. Chou, S. Yeung, N. Naik, A. Madani, A. Mottaghi, and J. Dean, "Deep learning-enabled medical computer vision," *Nature Medicine*, vol. 29, no. 2, pp. 113–124, 2023.
- [12] F. Jiang, Y. Jiang, H. Zhi, Y. Dong, H. Li, S. Ma, and Q. Wang, "Artificial intelligence in healthcare: Past, present and future," *Stroke and Vascular Neurology*, vol. 2, no. 4, pp. 230–243, 2017.

- [13] A. Rajkomar, J. Dean, and I. S. Kohane, "Machine learning in medicine," *New England Journal of Medicine*, vol. 380, no. 14, pp. 1347–1358, 2019.
- [14] B. Hendriks, A. van Engelen, and M. Hilgsmann, "Ethical implications of AI-driven clinical decision support systems on resource allocation," *BMC Medical Ethics*, vol. 25, no. 1, 115, 2024.
- [15] A. Holzinger, P. Kieseberg, E. Weippl, and A. M. Tjoa, "Causability and explainability of artificial intelligence in medicine," *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, vol. 9, no. 4, e1312, 2019.
- [16] M. A. Razzak, S. Naz, and A. Zaib, "Deep learning for medical image processing: Overview, challenges and the future," *Classification in BioApps*, Springer, pp. 323–350, 2018.
- [17] E. Tjoa and C. Guan, "A survey on explainable artificial intelligence (XAI): Toward medical XAI," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 32, no. 11, pp. 4793–4813, 2021.
- [18] H. Hildt, "What Is the Role of Explainability in Medical Artificial Intelligence? A Case-Based Approach," *Bioengineering (Basel)*, vol. 12, no. 4, 375, 2025.
- [19] M. Alizadehsani, R. Alizadeh, F. Shahraki, L. Benyamin, and S. Moosaei, "A review of Explainable Artificial Intelligence in healthcare," *Computers in Biology and Medicine*, vol. 164, 107408, 2024.
- [20] H. Chen, J. Zhang, and W. Xu, "Artificial intelligence in clinical decision support: Challenges and opportunities," *IEEE Journal of Biomedical and Health Informatics*, vol. 27, no. 8, pp. 4102–4113, 2023.
- [21] F. Mohsen, H. Ali, N. El Hajj, and Z. Shah, "Artificial intelligence-based methods for fusion of electronic health records and imaging data," *Scientific Reports*, vol. 12, article no. 17981, 2022.
- [22] S. Tian, N. An, and H. Mao, "Smart healthcare: making medical care more intelligent," *Urban Studies and Social Science*, vol. 16, no. 1, pp. 1–7, 2019.
- [23] B. Ç. Uslu and C. Benli, "Analysis of factors affecting IoT-based smart hospital design," *Journal of Cloud Computing: Advances, Systems and Applications*, vol. 9, article no. 30, 2020.
- [24] S. Topol, "High-performance medicine: The convergence of human and artificial intelligence," *Nature Medicine*, vol. 25, pp. 44–56, 2019.
- [25] A. Panch, H. Mattie, and R. Atun, "Artificial intelligence and algorithmic bias: Implications for health systems," *The Lancet Digital Health*, vol. 1, no. 5, pp. e200–e205, 2019.
- [26] Sabin, O. Information Sources in Romanian on Medicines. *Medic.ro* 2019, 4, 28. [Google Scholar] [CrossRef]
- [27] Shen, Z.; Spruit, M. Automatic Extraction of Adverse Drug Reactions from Summary of Product Characteristics. *Appl. Sci.* 2021, 11, 2663. [Google Scholar] [CrossRef]