

Leveraging Agentic AI To Optimize Hospital Operations and Clinical Decision Making

1st Allakonda Nichitha

Information Technology

Chaitanya Bharathi Institute of Technology

Hyderabad, India

anichitha2004@gmail.com

2nd Gadagotti Harika

Information Technology

Chaitanya Bharathi Institute of Technology

Hyderabad, India

harikagadagotti@gmail.com

3rd V.Deekshita

Information Technology

Chaitanya Bharathi Institute of Technology

Hyderabad, India

deekshitareddyvallur@gmail.com

4th B.Veera Jyothi

Assistant Professor, Department of Information Technology,

Chaitanya Bharathi Institute of Technology

Hyderabad, India

veerajyothi_it@cbi.ac.in

Abstract—Healthcare systems often face challenges such as delays in identifying patient deterioration, rising administrative workload, and inefficient use of hospital resources. While machine learning models are widely used to predict patient outcomes, they are typically deployed as standalone tools and are not well integrated into everyday hospital workflows.

This work presents a healthcare decision support system built using an agent-based artificial intelligence approach. It combines predictive analysis, retrieval of clinical knowledge, and operational decision-making within a multi-agent framework.

Patient risk is assessed using an XGBoost model based on physiological data. Relevant clinical guidelines are retrieved using similarity-based methods, while hospital resource allocation decisions are guided by a reinforcement learning model. The outputs from these components are then integrated to generate a unified decision for patient care and resource management.

The model shows high predictive performance with an accuracy of 0.9984, precision of 0.9953, recall of 1.000, F1-score of 0.9977, and ROC-AUC of 1.000. The system is implemented through a Streamlit interface that allows real-time input of patient data and generation of recommendations.

Index Terms—Agentic AI, Multi-Agent Systems, Clinical Decision Support, Healthcare Automation, Hospital Operations, MIMIC-III, Explainable AI.

I. INTRODUCTION

Healthcare systems today handle large numbers of patients while also managing administrative work and limited medical resources. Hospitals must observe patient conditions continuously and organize the availability of medical staff, equipment, and intensive care unit (ICU) facilities. When changes in a patient's condition are not identified early, medical intervention may be delayed, which can affect treatment results and patient safety [10][14].

Artificial Intelligence (AI) is increasingly used in healthcare environments for tasks such as disease diagnosis, patient risk prediction, medical image analysis, and monitoring of clinical information [8][10]. Machine learning methods allow large

volumes of clinical data to be examined and help reveal patterns related to patient outcomes. However, many predictive models are used as separate analytical tools. In practice they often produce predictions without being directly connected to hospital workflows or operational management systems [3][6].

Recent research has introduced Agentic Artificial Intelligence, where several autonomous agents cooperate to perform tasks in complex and changing environments [1][2]. In healthcare settings, this type of architecture can link predictive models with clinical knowledge sources and hospital management processes. Such coordination allows patient monitoring, decision support, and resource management to operate within the same computational system.

This paper introduces a healthcare decision support framework built on an agent-based AI architecture. The system brings together predictive analysis, clinical knowledge retrieval, and hospital operations planning. Its goal is to assess patient risk in real time while helping hospitals allocate resources effectively and support day-to-day clinical activities.

II. LITERATURE SURVEY

Artificial Intelligence (AI) is increasingly used in healthcare to support clinical decision making, patient monitoring, and hospital management. The growth of electronic health records (EHR) and large medical datasets has enabled the development of machine learning models capable of identifying complex relationships in clinical data. These models are commonly applied for disease prediction and patient risk assessment. Previous studies report that AI-based analytical methods can improve diagnostic accuracy and assist clinicians in predicting patient outcomes more effectively than traditional statistical approaches [10][14].

Machine learning techniques have also been investigated for early detection of patient deterioration. Predictive models trained on physiological measurements and clinical records are able to identify high-risk patients and support earlier medical

Identify applicable funding agency here. If none, delete this.

intervention. Wubineh et al. [3] reviewed several AI applications in clinical environments and emphasized the importance of predictive analytics in patient monitoring. Sahithya et al. [8] showed that techniques such as convolutional neural networks and natural language processing can be applied to medical images and electronic health records to improve diagnostic support.

Beyond predictive analysis, AI has also been applied to improve hospital operations. Healthcare organizations often deal with issues such as staff scheduling, resource allocation, and managing patient appointments. Knight et al. [9] studied AI-based scheduling systems in hospitals and found that they improved appointment handling and overall resource use. These findings indicate that intelligent optimization methods can help make healthcare workflows more efficient.

More recent work has introduced the idea of Agentic Artificial Intelligence, which focuses on autonomous systems that can reason, plan, and interact with other agents to handle complex tasks. Hinostroza Fuentes et al. [1] described AI systems with such capabilities and highlighted how agent-based approaches can enhance adaptability and decision-making in healthcare settings. In these systems, agents are able to process contextual information, communicate with one another, and coordinate their actions effectively.

In their overview of agentic AI architecture, Acharya et al. [2] examined the possibility of developing goal-oriented intelligent systems by integrating symbolic reasoning and reinforcement learning. These intelligent systems have the capacity to manage different tasks and adapt to different situations. These are important factors in healthcare environments, where clinical decisions are often subject to both the health conditions of the patients and the resource conditions of the hospital.

Large language models and generative AI systems have also been explored for healthcare-related decisions. Reddy, in his overview of the application of generative AI in medical systems, emphasized the issues of compliance, deployment, and governance of generative AI systems in healthcare environments. The clinical suggestions generated by AI systems are useful to physicians as they provide medical information that is context-sensitive, as evaluated by Liu et al. in their evaluation of the reliability of AI systems, although there are issues of reliability and inaccuracies in the results generated by the systems.

Ethical considerations are yet another facet of the implementation of AI systems. Bekbolatova et al. [5] identified the challenges of algorithmic bias, transparency, and accountability of AI systems while working in a healthcare environment. Esmailzadeh [6] also identified organizational, regulatory, and infrastructural limitations of AI systems while working in a healthcare environment.

Economic and policy research has also identified inefficiencies while working within a healthcare environment. Papanicolas et al. [13] and Kyle and Frakt [12] identified the inefficiencies of workload and fragmented work management as major factors contributing to the cost of healthcare. Cutler

and Ly [16] identified the cost of administration as a major component of healthcare expenditure. This clearly indicates the scope of intelligent automation in this regard.

Despite the research and developments carried out on healthcare-related AI systems, several limitations have been identified. The limitations include the focus of several AI systems on performing a particular predictive task without taking into account the incorporation of such predictions within the larger hospital environment.

To address these issues, this work proposes an Agentic AI-based multi-agent framework that combines predictive analytics, knowledge retrieval, and operational decision support. The framework enables collaboration between predictive models and intelligent agents in order to support clinical decision making and improve hospital operational efficiency.

III. PROPOSED METHODOLOGY

The proposed system introduces an Agentic Artificial Intelligence-based multi-agent framework designed to support healthcare decision making and hospital resource management.

Agentic AI systems are capable of autonomous reasoning, collaborative problem solving, and adaptive decision making in complex environments such as healthcare systems [1], [2].

The framework integrates predictive analytics, clinical knowledge retrieval, and operational optimization within a coordinated agent architecture. Artificial intelligence techniques have increasingly been applied in healthcare environments to improve clinical decision support, patient monitoring, and hospital workflow management [10], [14].

Integration of these components enables the system to process physiological patient data, estimate mortality risk, and generate both clinical and operational recommendations. AI-based analytical systems can support clinicians in recognizing high-risk patient conditions at an early stage and facilitate timely medical intervention, which contributes to improved patient safety and treatment outcomes [3], [8], [10].

The system architecture adopts a multi-agent structure coordinated through a central orchestration mechanism. Each intelligent agent is designed to handle a particular task and communicates with other components through an orchestrator. Agent-based and agentic AI systems support independent coordination among modules, which makes them suitable for healthcare settings where predictive analysis and hospital operations need to work in alignment [1], [2].

A. Data Acquisition and Preprocessing

The system uses clinical data from the MIMIC-III electronic health records dataset, which includes detailed information on patients admitted to intensive care units. Such electronic health records are commonly used in healthcare research to build machine learning models for predicting patient risk and supporting clinical decisions [8], [10].

From this dataset, important physiological features are chosen to develop the prediction model. Well-structured medical data helps AI techniques analyze a patient's condition and

detect potential health risks based on patterns in the data [14]

In the preprocessing phase, vital signs are taken from the CHARTEVENTS table, with attention given to key physiological parameters. This step is important in healthcare analytics because it prepares the data by cleaning and organizing it, making it suitable for training machine learning models [3].

- Heart Rate (HR)
- Oxygen Saturation (SpO₂)
- Body Temperature (Temp)

These measurements are then filtered in accordance with certain clinical item identifiers and aggregated in the form of mean statistics. The vital sign information, which is extracted from the database, is then structured in a suitable form using pivot operations, which arrange the patient records according to the subject identifier and the timestamp. This type of preprocessing technique is often used in the context of machine learning-based healthcare systems, where the EHR information needs to be preprocessed for analysis [8], [10].

Next, the vital sign information, which has been preprocessed, is joined with the ADMISSIONS table. In this case, records with incomplete information are removed from the dataset, which then becomes the input for the predictive model, thereby facilitating the analysis of the health conditions of the patients [14].

B. Clinical Agent – Mortality Risk Prediction

The Clinical Agent performs patient risk prediction using machine learning methods. In this, an XGBoost classifier is selected because its ability to handle nonlinear relationships and heterogeneous healthcare datasets effectively [8], [10].

The predictive model is trained using physiological parameters extracted from patient records. Machine learning models trained on electronic health records have demonstrated strong capability in predicting patient outcomes and supporting clinical decision-making in healthcare environments [3], [14].

The parameters used in the model include:

- Age
- Heart rate
- Respiratory rate
- Blood pressure
- Body temperature
- Oxygen saturation (SpO₂)
- Glucose level
- White blood cell count
- Creatinine level
- ICU stay duration

The trained model produces a probability score that represents the likelihood of patient mortality risk. This probability value is used to classify patients into risk categories based on a predefined threshold. Patients with higher predicted probabilities are categorized as high-risk cases, while lower probabilities indicate a relatively stable condition.

The generated risk score supports healthcare professionals in identifying patients who may require closer observation

or early clinical intervention. By providing a quantitative estimation of patient risk, the Clinical Agent contributes to timely monitoring and assists clinicians in making informed treatment decisions [3], [10].

C. Admin Agent – Retrieval Augmented Guidelines

The Administrative Agent retrieves appropriate clinical guidelines based on the predicted patient risk level. To achieve this, the system uses Sentence Transformer embeddings combined with FAISS vector similarity search. Retrieval-based AI systems are increasingly used in healthcare applications to provide context-aware clinical recommendations and assist medical professionals in decision making [7], [11].

A set of predefined medical guidelines is encoded into vector representations using a transformer-based language model. These embeddings are indexed using FAISS to enable efficient semantic search. Such embedding-based retrieval techniques are widely applied in healthcare AI systems to access relevant medical knowledge from large information sources [8], [10].

When a patient risk level is determined by the Clinical Agent, the Administrative Agent generates a query representing the patient condition. The system then retrieves the most relevant guideline from the indexed knowledge base. This mechanism allows the system to provide context-aware medical recommendations aligned with patient status, thereby supporting clinicians in selecting appropriate treatment protocols [7], [11].

D. Operations Agent – Resource Optimization

The Operations Agent focuses on hospital resource management using reinforcement learning techniques. AI-based optimization methods have increasingly been explored to improve hospital workflow efficiency and resource utilization in healthcare systems [9], [16]. A simulated hospital environment is developed using the OpenAI Gym framework to represent hospital resource conditions.

The environment models the hospital state using three variables:

- Available ICU beds
- Available medical staff
- Incoming patient count

A Proximal Policy Optimization (PPO) reinforcement learning algorithm is used to train the agent to learn optimal resource management actions. Reinforcement learning techniques enable intelligent agents to learn adaptive decision-making strategies in dynamic environments, making them suitable for complex healthcare operational scenarios [2].

Possible actions include:

- Allocating ICU beds
- Calling additional medical staff
- Maintaining current resource allocation

The trained agent evaluates the hospital environment and selects the action that maximizes resource utilization efficiency. Such AI-driven operational optimization can reduce administrative burden and improve the efficiency of healthcare systems [12], [16].

E. Orchestrator Agent – Multi-Agent Coordination

The Orchestrator Agent serves as the central coordination component of the system. Multi-agent architectures enable different intelligent components to collaborate and perform coordinated decision making in complex environments such as healthcare systems [1], [2]. In the proposed framework, the orchestrator integrates outputs from the Clinical Agent, Administrative Agent, and Operations Agent to generate a unified decision.

The orchestrator receives patient physiological inputs and forwards them to the Clinical Agent for risk prediction. Based on the predicted risk level, the Administrative Agent retrieves the most relevant clinical guideline. Simultaneously, the Operations Agent analyzes the hospital environment and recommends an appropriate resource management action. AI-driven systems can assist clinicians by coordinating predictive insights with operational decision support, thereby improving healthcare workflow efficiency [3], [10].

The orchestrator aggregates these outputs into a structured response that includes the predicted risk level, mortality probability, clinical guideline, and recommended operational action. Such coordinated agent-based systems enable more adaptive and context-aware healthcare decision support frameworks [1], [7].

F. Interactive Dashboard

To enable practical use in clinical environments, the proposed system is integrated into an interactive dashboard developed using Streamlit. Interactive decision-support systems allow healthcare professionals to analyze patient data and receive AI-assisted recommendations in real time, improving clinical workflow efficiency and decision making [7], [10].

The dashboard allows healthcare professionals to input patient parameters and obtain real-time predictions and recommendations. AI-powered healthcare interfaces are increasingly used to support clinicians by providing visual representations of model predictions and clinical insights derived from patient data [8], [11].

The interface displays:

- Patient risk probability
- Model confidence score
- Feature importance explanations
- Clinical guideline recommendations
- Resource allocation suggestions

This visualization platform helps clinicians quickly interpret model outputs and make informed decisions regarding patient care, thereby improving the usability and practical applicability of AI-based healthcare systems [7], [11].

G. Final Evaluation and Reporting

To enable real-world usability, the system is deployed through an interactive Streamlit dashboard. Interactive AI-based decision support systems allow healthcare professionals to analyze patient data and receive real-time predictions that assist in clinical decision making [7], [10].

The dashboard allows healthcare professionals to input patient vital signs and obtain real-time predictions. AI-driven healthcare interfaces have been widely used to visualize predictive outputs and support clinicians in evaluating patient conditions using clinical data [8], [11].

Users can provide the following inputs:

- Heart rate
- Oxygen saturation
- Body temperature

After submission, the system displays the predicted risk level, mortality probability, and recommended actions. This interface enables clinicians to quickly interpret results and make informed medical decisions, improving the practical usability of AI-based healthcare systems [7], [11].

IV. SYSTEM ARCHITECTURE

The architecture of the proposed Agentic AI-based Healthcare Decision Support System is illustrated in Fig. 1. The system integrates machine learning, retrieval-based knowledge systems, and reinforcement learning within a multi-agent framework to support clinical decision-making and hospital resource management. Agentic AI architectures enable multiple intelligent components to collaborate and perform coordinated decision-making in complex environments such as healthcare systems [1], [2].

The process starts with the data acquisition stage, where patient clinical data are obtained from the MIMIC-III dataset. This dataset contains electronic health records including physiological measurements and patient outcomes. Electronic health record datasets are widely used in healthcare research to develop machine learning models for patient risk prediction and clinical decision support [8], [10]. Relevant features such as heart rate, oxygen saturation, and body temperature are extracted from the raw data.

In the next stage, data preprocessing and feature engineering are performed to prepare the dataset for model training. This step involves filtering relevant clinical variables, handling missing values, and transforming the data into a structured format suitable for machine learning algorithms. Data preprocessing is a critical step in healthcare analytics to ensure reliable and accurate model training using clinical datasets [3], [14].

The processed data are then used to train an XGBoost-based predictive model, which is responsible for estimating the probability of patient mortality risk. Machine learning models trained on clinical data have demonstrated strong capability in predicting patient outcomes and assisting healthcare professionals in decision-making processes [8], [10].

The predictive model is integrated into a Streamlit-based dashboard, which acts as the main user interface and multi-agent controller. Healthcare professionals can enter patient vital signs through the dashboard to obtain risk predictions and clinical recommendations. Interactive AI-based decision-support systems enable clinicians to interpret model outputs more effectively and support real-time healthcare decision making [7], [11].




Fig. 1. Architecture of the proposed Agentic AI-based healthcare decision support system.

The system then activates three specialized intelligent agents:

A. Clinical Agent

The Clinical Agent performs patient risk prediction using the trained machine learning model. Machine learning techniques applied to clinical datasets enable accurate prediction of patient health risks and support early medical intervention in healthcare environments [8], [10].

B. Administrative Agent

The Administrative Agent retrieves appropriate clinical protocols using retrieval-augmented generation (RAG) techniques and semantic similarity search. AI-based clinical knowledge retrieval systems assist healthcare professionals by providing context-aware medical recommendations derived from clinical data and medical guidelines [7], [11].

C. Operations Agent

The Operations Agent optimizes hospital resource allocation using reinforcement learning (RL) to recommend actions such as ICU bed allocation or staff adjustments. AI-driven operational optimization techniques have been used to improve hospital workflow efficiency and resource utilization in healthcare systems [9], [16].

These agents communicate through an Orchestrator Agent, which acts as the central coordination mechanism of the system. Multi-agent architectures enable intelligent components to collaborate and perform coordinated decision making in complex environments such as healthcare systems [1], [2]. The orchestrator integrates predictions from the Clinical

Agent, clinical guidelines from the Administrative Agent, and operational recommendations from the Operations Agent.

Finally, the combined output is presented through the Risk Dashboard Output, which provides healthcare professionals with a comprehensive summary of the patient's condition, recommended clinical protocol, and operational action plan. Such AI-based decision support systems improve clinical workflow efficiency and support informed healthcare decision making [7], [10].

Advantages:

The proposed Agentic AI-based multi-agent framework provides several advantages over traditional hospital AI systems.

- **Autonomous and Context-Aware Decision-Making:** Each agent—Clinical, Administrative, and Operational—functions independently while remaining connected through a central orchestrator. This structure enables real-time responses and allows the system to adjust to changing clinical and operational conditions, resulting in more context-aware and timely decisions.
- **Enhanced Efficiency and Predictive Accuracy:** By integrating predictive analytics with reinforcement learning (RL), the framework supports better clinical and administrative outcomes. It helps in early identification of patient risks, reduces response time, and improves the use of hospital staff and resources.
- **Transparency and Explainability:** Using SHAP-based methods, the system presents each prediction along with the factors that influenced it. This improves clarity for clinicians, builds trust in the results, and supports more responsible use of AI in healthcare.
- **Cross-Functional Integration and Workflow Automation:** The system connects areas that are usually handled separately, such as patient monitoring, hospital logistics, and administrative tasks. By coordinating these activities, it helps daily operations run more smoothly, improves communication between teams, and reduces manual work, creating a more streamlined healthcare workflow.
- **Scalability, Accessibility, and Open-Source Flexibility:** Built with accessible, open-source tools like Google Colab, Streamlit, and Docker, the system can be scaled as needed. This makes it practical for both large hospitals and smaller regional healthcare centers.
- **Ethical Governance and Human-Centered Design:** The framework includes built-in safeguards and keeps clinicians involved in the decision-making process. This ensures patient safety, supports ethical practices, and maintains accountability with healthcare professionals.

V. SYSTEM IMPLEMENTATION AND OUTCOMES

The proposed Agentic AI framework was developed in Python using a combination of machine learning and AI libraries. The implementation integrates predictive modeling, semantic retrieval, reinforcement learning, and an

interactive web interface to provide a complete healthcare decision support system.

A. System Implementation

The proposed Agentic AI-based healthcare decision support system was implemented using Python and several machine learning libraries. The system integrates predictive modeling, clinical guideline retrieval, reinforcement learning-based resource optimization, and an interactive web interface for real-time decision support. Artificial intelligence techniques have increasingly been applied in healthcare systems to support clinical decision making and improve hospital workflow efficiency [3], [10].

The predictive component of the system was developed using the XGBoost classification algorithm, which was trained on patient physiological features extracted from the dataset. Machine learning models trained on electronic health records have demonstrated strong capability in predicting patient outcomes and supporting healthcare analytics [8], [14]. The model predicts patient risk levels based on parameters such as heart rate, oxygen saturation, body temperature, and other clinical indicators.

The system architecture follows a multi-agent framework, where individual agents perform specialized tasks. Agent-based architectures enable different intelligent components to collaborate and perform coordinated decision making in complex healthcare environments [1], [2]. The Clinical Agent performs risk prediction using the trained XGBoost model. The Administrative Agent retrieves appropriate clinical guidelines based on predicted patient conditions using semantic retrieval techniques. The Operations Agent recommends hospital resource allocation strategies using reinforcement learning to optimize ICU bed usage and staff availability, which can improve hospital workflow efficiency and resource utilization [9], [16]. All agents communicate through an Orchestrator Agent, which combines the outputs of individual agents and generates the final coordinated decision. Such AI-driven decision support systems assist clinicians by integrating predictive insights with operational recommendations in real time [7], [11]. The entire system is deployed using a Streamlit-based dashboard, enabling healthcare professionals to enter patient parameters and obtain real-time clinical recommendations.

B. Model Training Results

The predictive model was trained using processed clinical data and evaluated using standard classification metrics such as accuracy, precision, recall, F1-score, and ROC-AUC. These metrics are widely used in healthcare applications to assess model performance and reliability.

The results demonstrate the strong predictive capability of the proposed XGBoost model. The model achieved an accuracy of 0.9984, precision of 0.9953, recall of 1.000, F1-score of 0.9977, and ROC-AUC score of 1.000. These values indicate that the model can effectively distinguish

between high-risk and low-risk patients with minimal misclassification. In particular, the recall value of 1.000 ensures that all critical patient cases are correctly identified, which is highly important in healthcare scenarios. The model training process and evaluation metrics obtained from the training stage are shown in Fig. 3.




Fig. 2. Model training results showing evaluation metrics including accuracy, precision, recall, F1-score, and confusion matrix for the XGBoost model.

The model training results, including evaluation metrics and confusion matrix, are illustrated in Fig. 2. The confusion matrix confirms that the majority of patient cases are correctly classified, with very few errors, demonstrating the robustness of the model.

Metric	Existing System (%)	Proposed System (%)
Accuracy	94.00	99.92
Precision	92.00	99.77
Recall	90.00	100.00
F1-Score	91.00	99.88
F1-Score	91.00	99.88

Fig. 3. Performance Comparison of Existing and Proposed System

Furthermore, Table 1 presents a comparative analysis between the existing system and the proposed system. The results clearly show that the proposed system significantly outperforms the existing approach across all performance metrics. The accuracy improves from 94.00% to 99.92%, while precision increases from 92.00% to 99.77%. The recall reaches 100.00%, ensuring complete identification of high-risk patients. Similarly, the F1-score improves from 91.00% to 99.88%, indicating a balanced enhancement in both precision and recall.

The confusion matrix shown in Fig. 3 confirms that the model correctly classified most patient records with only a small number of misclassifications.

C. Clinical Parameter Input Interface

The developed system provides a user-friendly interface for healthcare professionals to enter patient clinical parameters. Interactive healthcare dashboards allow clinicians to analyze patient data and obtain AI-assisted predictions for clinical decision making [7], [11].



Fig. 4. Streamlit interface used to enter patient clinical parameters for risk assessment.

Fig. 4 shows the input interface of the proposed system. Healthcare staff can enter patient physiological values through this dashboard, and the system processes the input to generate risk predictions.

D. Low-Risk Patient Prediction Output

When patient parameters fall within normal clinical ranges, the Clinical Agent predicts a low-risk condition. In such cases, the Administrative Agent retrieves appropriate medical guidelines recommending routine monitoring, which assists clinicians in managing patient conditions efficiently [7], [10].




Fig. 5. Dashboard output showing prediction results for a low-risk patient including mortality probability and clinical recommendation.

As shown in Fig. 5, the system predicted a low mortality probability of 0.003 with a model confidence of 0.997. Based on this prediction, the system recommends routine patient monitoring and standard ward care.

E. High-Risk Patient Prediction Output

When abnormal clinical parameters are detected, the system identifies the patient as high risk and triggers emergency recommendations. AI-based predictive models can assist clinicians in detecting critical patient conditions at an early stage and enabling timely medical intervention [3], [8].



Fig. 6. Dashboard output showing prediction results for a high-risk patient including mortality probability and emergency clinical recommendations.

Fig. 6 illustrates the high-risk scenario where the model predicted a mortality probability of 0.998. In this case, the Administrative Agent recommended initiating a sepsis protocol and administering antibiotics immediately.

F. Resource Allocation Recommendation

The Operations Agent evaluates hospital resources and suggests operational actions such as ICU admission or staff allocation. AI-based hospital management systems have been shown to improve resource utilization and reduce operational inefficiencies in healthcare institutions [9], [16].

As shown in Fig. ??, the system recommends immediate ICU admission and continuous monitoring when a patient is classified as high risk.

G. Final Orchestrated Decision

The Orchestrator Agent combines the outputs from the Clinical Agent, Administrative Agent, and Operations Agent to produce the final decision summary. Multi-agent AI systems enable coordinated decision making by integrating predictive analytics with operational insights [1], [2].

The integrated decision output ensures that clinical prediction, treatment guidelines, and operational resource management are considered simultaneously, thereby improving the efficiency and reliability of healthcare decision-support systems [7], [10].

VI. CONCLUSION

The proposed Agentic AI framework takes a practical and forward-looking approach to improving healthcare by combining smart automation with advanced reasoning and coordinated agent interactions. Using predictive analytics, retrieval-based medical recommendations, and

reinforcement learning for operational improvement, the system supports real-time clinical decisions and streamlines hospital workflows [1][2][3].

Built on a multi-agent setup, each agent focuses on a specific area clinical prediction, administrative support, operations management, or patient engagement [2][4]. The Clinical Agent identifies early signs of patient deterioration, while the Administrative and Operations Agents handle tasks like claims processing, bed allocation, and scheduling to improve efficiency [1][7][11][5][6][9][12]. Coordinated through a central orchestrator, these agents support a smooth workflow and help clinicians respond faster with better-informed decisions [2][8].

A key strength of the system is the use of explainable AI with SHAP, which makes the model's decisions easier to understand and helps build trust among clinicians [1][2][5]. This level of transparency also supports the ethical use of AI in healthcare [1][5][7]. As the agents work together, the system learns from real-world data, gradually improving and connecting clinical insights with hospital operations [3][9][12][13].

In summary, the Agentic AI framework moves toward a more reliable and patient-focused approach to healthcare. By combining real-time monitoring, coordinated agent actions, and human involvement, it provides a strong base for systems that can improve the quality and accuracy of clinical decisions [1][2][5][14].

Future Work:

To further improve and expand this framework, the following research and development directions are suggested:

- Increase Medical Knowledge Bases: Enrich the Admin Agent with access to current clinical guidelines, new research and hospital-specific protocols.
- Implementation in Real-Life Hospital setting: Hold preliminary studies in real clinical settings to determine how the entire system can fit in current EHR based systems and whether it enhances day-to-day operations.
- Sensor Data Live IoT and Wearable Data: Incorporate patient continuous health data (wearable sensors and IoT devices) to be able to detect and respond more quickly to changes in the patients.
- Federated and Privacy-Preserving Learning: Implement federated learning in a way that will allow hospitals to jointly train models without transferring sensitive patient information, which will preserve privacy and security.
- Generative and Conversational AI Incorporation: Add conversational agents to assist the communication with patients and apply generative reasoning models to offer more adaptive and context-sensitive clinical recommendations.
- Continuous Self-Learning and Adaptation: Design re-

inforcement learning-based feedback systems to enable autonomous agents to continuously learn how to improve their decision making strategies using real world outcomes.

REFERENCES

- [1] V. G. Hinostroza Fuentes, H. Abdul Karim, M. J. Toledo Tan, and N. AlDahoul, "Ai with agency: a vision for adaptive, efficient, and ethical healthcare," *Frontiers in Digital Health*, vol. 7, p. 1600216, May 2025.
- [2] D. B. Acharya, K. Kuppan, and B. Divya, "Agentic ai: autonomous intelligence for complex goals—a comprehensive survey," *IEEE Access*, vol. 13, February 2025.
- [3] B. Z. Wubineh, F. G. Deriba, and M. M. Woldeyohannis, "Exploring the opportunities and challenges of implementing artificial intelligence in healthcare: a systematic literature review," *Urologic Oncology: Seminars and Original Investigations*, vol. 42, no. 1, January 2024.
- [4] S. Khaleelullah, H. Kumar, G. Rahul, R. Naik, and S. Teja, "Dietrx: machine learning enhanced disease specific nutrition and precautions," in *Proceedings of ICISC 2024*. IEEE, January 2024.
- [5] M. Bekbolatova, J. Mayer, C. W. Ong, and M. Toma, "Transformative potential of ai in healthcare: definitions, applications, and navigating the ethical landscape and public perspectives," *Healthcare*, vol. 12, no. 2, February 2024.
- [6] P. Esmailzadeh, "Challenges and strategies for wide-scale ai deployment in healthcare practices: a perspective for healthcare organizations," *Artificial Intelligence in Medicine*, vol. 151, March 2024.
- [7] S. Reddy, "Generative ai in healthcare: an implementation science informed translational path on application, integration and governance," *Implementation Science*, vol. 19, p. 27, February 2024.
- [8] B. Sahithya, M. S. Prasad, K. M. Krishna, A. C. Devarlla, and T. R. Yashavanth, "Empowering healthcare with ai: advancements in medical image analysis, electronic health records analysis, and ai-driven chatbots," in *Proceedings of INOCON 2024*. IEEE, January 2024.
- [9] D. R. T. Knight, C. A. Aakre, C. V. Anstine, B. Munipalli, P. Biazar, G. Mitri *et al.*, "Artificial intelligence for patient scheduling in the real-world health care setting: a metanarrative review," *Health Policy and Technology*, vol. 12, no. 4, October 2023.
- [10] S. A. Alowais, S. S. Alghamdi, N. Alsuhbany *et al.*, "Revolutionizing healthcare: the role of artificial intelligence in clinical practice," *BMC Medical Education*, vol. 23, no. 1, December 2023.
- [11] S. Liu, A. P. Wright, B. L. Patterson, J. P. Wanderer, S. D. Nelson *et al.*, "Using ai-generated suggestions from chatgpt to optimize clinical decision support," *Journal of the American Medical Informatics Association*, vol. 30, no. 7, July 2023.
- [12] M. A. Kyle and A. B. Frakt, "Patient administrative burden in the us health care system," *Health Services Research*, vol. 56, no. Suppl. 1, August 2021.
- [13] I. Papanicolas, L. R. Woskie, D. Orlander, E. J. Orav, and A. K. Jha, "The relationship between health spending and social spending in high-income countries: how does the us compare?" *Health Affairs*, vol. 38, no. 9, September 2019.
- [14] F. Jiang, Y. Jiang, H. Zhi, Y. Dong, H. Li, S. Ma *et al.*, "Artificial intelligence in healthcare: past, present and future," *Stroke and Vascular Neurology*, vol. 2, no. 4, December 2017.
- [15] Y. Zhang, R. Padman, and J. E. Levin, "Paving the cowpath: data-driven design of pediatric order sets," *Journal of the American Medical Informatics Association*, vol. 21, no. 2, March 2014.
- [16] D. M. Cutler and D. P. Ly, "The (paper) work of medicine: understanding international medical costs," *Journal of Economic Perspectives*, vol. 25, 2011.