

Vol 2 Issue 02 (2026)

<http://fsrma.org.hostajournal.com/>

MedAgentOps: A Distributed Multi-Agent Collaborative Reasoning Infrastructure Based on Memory-Augmented Graph Reasoning Network for Emergency Department Flow Forecasting and Resource Optimization

Hao Hu

Bowling Green State University, Bowling Green, OH, USA, 43403
felixhuhao@gmail.com

Yan Zeng

University of California, Berkeley, CA, USA, 94720
yan.zeng@berkeley.edu

Weijia Li

Shandong University, Jinan, China
1765368138@qq.com

Abstract.

Emergency departments (EDs) in modern hospitals face persistent challenges including unpredictable patient influx, resource scarcity, and inefficient operational coordination, which significantly degrade service quality and response efficiency. To address these issues, this study proposes MedAgentOps, a distributed multi-agent collaborative reasoning infrastructure built upon a Memory-Augmented Graph Reasoning Network (MAGR-Net) for emergency flow forecasting and resource optimization. The proposed framework integrates four key components: Dynamic Temporal Graph Fusion (DTGF) for modeling heterogeneous temporal dependencies in emergency-related factors, a Shared Medical Memory (SMM) module for consolidating short-term operational states and long-term historical knowledge, an Agent Graph Reasoning Module (AGRM) for adaptive coordination among forecasting, scheduling, and resource allocation agents, and a Multi-Objective Proximal Policy Optimization (MOPPO) strategy for optimizing operational efficiency under multiple constraints. These modules jointly enable end-to-end reasoning from patient arrival prediction to real-time resource scheduling in a unified multi-agent system. Experimental results on a real-world emergency department dataset show that the proposed framework achieves a Mean Absolute Percentage Error (MAPE) of 8.73%, outperforming all baseline models significantly. In addition, MedAgentOps reduces average patient waiting time by approximately 17.9% compared with the strongest baseline (GNN-RL) while improving resource utilization to 91.3%. These results confirm the effectiveness of integrating multi-agent reasoning with memory-augmented graph learning for intelligent healthcare optimization.

Keywords:

Multi-Agent Systems;
Emergency Department Forecasting;
Distributed Artificial Intelligence Infrastructure;
Collaborative Reasoning;
Resource Optimization;
Graph Neural Networks;
Reinforcement Learning;
Healthcare Informatics

1. INTRODUCTION

Emergency Departments (EDs) are critical components of modern healthcare systems, responsible for handling unpredictable patient arrivals and providing timely medical services under severe resource constraints. In recent years, increasing population density, seasonal epidemics, and public health emergencies have significantly intensified the pressure on ED operations, leading to overcrowding, prolonged waiting times, and inefficient utilization of medical resources. Consequently, accurate emergency flow forecasting and efficient resource optimization have become essential for improving hospital operational efficiency and patient outcomes. Although recent advances in time-series forecasting models such as Long Short-Term Memory (LSTM), Transformer, and Temporal Fusion Transformer (TFT) have achieved promising results in predicting patient arrival patterns, these methods largely focus on isolated prediction tasks and fail to capture complex operational dependencies within hospital systems. Moreover, existing approaches rarely incorporate decision-making mechanisms or collaborative reasoning capabilities required for real-world emergency resource scheduling. To address these limitations, this study proposes MedAgentOps, a distributed multi-agent collaborative reasoning infrastructure designed for emergency department flow forecasting and resource optimization. At its core, MedAgentOps is powered by a novel Memory-Augmented Graph Reasoning Network (MAGR-Net), which enables coordinated intelligence across multiple specialized agents, including prediction, scheduling, resource allocation, and medical knowledge reasoning agents. The framework integrates dynamic temporal modeling, shared memory mechanisms, and graph-based agent collaboration into a unified system for end-to-end healthcare operation optimization. Specifically, MedAgentOps incorporates four key components: (1) a Dynamic Temporal Graph Fusion (DTGF) module for modeling heterogeneous temporal dependencies among clinical and operational factors; (2) a Shared Medical Memory (SMM) mechanism for preserving both short-term operational states and long-term historical knowledge; (3) an Agent Graph Reasoning Module (AGRM) for adaptive coordination and information propagation among heterogeneous agents; and (4) a Multi-Objective Proximal Policy Optimization (MOPPO) strategy for optimizing resource allocation under multiple operational constraints.

The main contributions of this work are summarized as follows:

- (1) We propose MedAgentOps, a distributed multi-agent infrastructure for joint emergency flow forecasting and resource optimization.
- (2) We design MAGR-Net, a memory-augmented graph reasoning model enabling collaborative decision-making across agents.
- (3) We introduce a shared medical memory mechanism to enhance long-term operational knowledge utilization.
- (4) We develop a graph-based agent reasoning module for dynamic inter-agent coordination.
- (5) We formulate a multi-objective reinforcement learning strategy to optimize ED resource efficiency and service quality.

2. Literature Review

Recent years have witnessed rapid advances in artificial intelligence (AI), particularly in deep learning, graph neural networks (GNNs), reinforcement learning (RL), and large language model (LLM)-based agent systems. These developments have significantly influenced intelligent healthcare systems, time-series forecasting, and distributed decision-making frameworks. In this section, we review related work in four key directions closely relevant to this study: healthcare time-series forecasting, reinforcement learning for resource optimization, graph-based reasoning models, and multi-agent AI systems for decision-making.

2.1 Healthcare Time-Series Forecasting for Emergency Department Management

Accurate prediction of emergency department (ED) patient arrivals is a fundamental task for hospital operation optimization. Early approaches relied on statistical models such as ARIMA and exponential smoothing, but their limited ability to capture nonlinear temporal dependencies restricts their performance. With the rise of deep learning, recurrent neural networks (RNNs) and Long Short-Term Memory (LSTM) models have become widely used in healthcare forecasting tasks [1]. More recently, attention-based architectures such as the Transformer model [2] and Temporal Fusion Transformer (TFT) [3] have demonstrated strong performance in modeling long-range temporal dependencies and multi-source medical data. Studies such as Ouyang et al. (2021) further show that transformer-based models outperform traditional RNNs in ED arrival prediction under complex seasonal patterns. However, most existing methods treat forecasting as an isolated task and ignore downstream operational decision-making requirements such as staffing and resource allocation. Furthermore, hybrid models integrating external factors such as weather, epidemics, and public events have been explored [4], but these approaches still lack integration with system-level optimization frameworks. This limitation motivates the need for unified architectures that jointly consider prediction and operational decision-making, as addressed in our MedAgentOps framework.

2.2 Reinforcement Learning for Healthcare Resource Optimization

Reinforcement learning (RL) has been widely applied to healthcare resource allocation problems, including patient scheduling, bed management, and ICU optimization. Mnih et al. introduced Deep Q-Network (DQN), which laid the foundation for RL-based decision systems [5]. Building on this, Proximal Policy Optimization (PPO) [6] has been widely adopted due to its stability and efficiency in continuous decision spaces. In healthcare applications, Yu et al. (2020) proposed an RL-based hospital bed allocation system that dynamically adjusts resource distribution based on patient inflow. Similarly, Li et al. applied multi-objective reinforcement learning to optimize emergency department waiting time and staffing costs simultaneously [7]. However, these methods typically rely on simplified state representations and lack integration with predictive models or structured reasoning modules. Recent studies have attempted to combine RL with forecasting models, but most frameworks remain loosely coupled, limiting their ability to perform end-to-end decision optimization in real-world hospital environments. In contrast, our proposed MOPPO module integrates prediction, reasoning, and optimization within a unified multi-agent framework.

2.3 Graph Neural Networks for Structured Reasoning and Healthcare Systems

Graph Neural Networks (GNNs) have emerged as a powerful tool for modeling relational structures in complex systems. Kipf and Welling introduced Graph Convolutional Networks (GCN) [8], which enabled learning on non-Euclidean data structures. Subsequently, Graph Attention Networks (GAT) [9] introduced attention mechanisms for adaptive neighbor aggregation. In healthcare, GNNs have been applied to disease prediction, patient similarity modeling, and medical knowledge graph reasoning. For example, Shang et al. (2019) used GNNs for patient trajectory prediction, demonstrating strong performance in capturing temporal and relational dependencies. However, most existing GNN-based healthcare models focus on static graph structures and do not consider dynamic operational environments such as emergency departments. More recently, dynamic graph learning methods such as Temporal Graph Networks (TGN) [10] have been proposed to model evolving interactions. Nevertheless, these approaches still lack integration with decision-making agents and memory mechanisms, which are critical for real-time hospital operation optimization.

2.4 Multi-Agent Systems and LLM-Based Collaborative Reasoning

Multi-agent systems (MAS) have been extensively studied for distributed decision-making and collaborative problem-solving. Early works focused on rule-based agent communication protocols and task allocation strategies. With the emergence of deep learning, multi-agent reinforcement learning (MARL) has become a dominant paradigm for cooperative systems [11]. Recently, LLM-based agents have gained significant attention due to their strong reasoning and tool-use capabilities. Frameworks such as AutoGPT and LangChain demonstrate how LLM agents can perform complex task decomposition and execution. However, most existing systems rely on sequential prompting strategies rather than structured collaboration mechanisms. Research on agent communication graphs and memory-augmented reasoning remains limited. Existing studies rarely integrate long-term memory, structured reasoning graphs, and reinforcement learning within a unified infrastructure. To address these gaps, our proposed MAGR-Net introduces a memory-augmented graph reasoning mechanism to enable structured collaboration among multiple specialized agents in healthcare environments.

2.5 Summary of Related Work

In summary, existing research has made significant progress in healthcare forecasting, reinforcement learning-based optimization, graph neural networks, and multi-agent systems. However, most methods remain fragmented: forecasting models are isolated from decision-making systems, RL-based approaches lack predictive awareness, and multi-agent systems often fail to incorporate structured memory and graph reasoning. These limitations highlight the necessity of a unified framework that integrates temporal forecasting, memory-augmented reasoning, and multi-agent collaboration, which motivates the proposed MedAgentOps and MAGR-Net architecture.

3. Methodology

3.1 Overview of MedAgentOps Framework

The proposed MedAgentOps is a distributed multi-agent collaborative reasoning infrastructure designed for emergency department flow forecasting and resource optimization. The core idea is to transform emergency department management into a unified reasoning process that integrates temporal prediction, knowledge retrieval, inter-agent collaboration, and optimization decision-making. Unlike conventional forecasting or scheduling systems that operate independently, MedAgentOps formulates the problem as a coupled spatio-temporal decision system driven by multiple interacting intelligent agents. At the heart of the framework lies the Memory-Augmented Graph Reasoning Network (MAGR-Net), which serves as the central reasoning engine for all agents. Given a sequence of multi-source emergency data $X_t = \{x_t^1, x_t^2, \dots, x_t^n\}$, including patient arrivals, resource status, and contextual signals, the system learns a unified representation:

$$H_t = F_{DTGF}(X_t)$$

where F_{DTGF} denotes the Dynamic Temporal Graph Fusion module that encodes heterogeneous temporal and structural dependencies. The resulting representation is then propagated through memory-enhanced and graph-based reasoning modules to generate forecasting and decision outputs.




Figure 1. Overall flowchart of the model.

3.2 Dynamic Temporal Graph Fusion (DTGF)

To capture complex dependencies among heterogeneous emergency department variables, drawing methodological parallels from dynamic temporal fusion approaches used in supply chain risk analysis [12], we model the system as a dynamic temporal graph $G_t = (V, E_t)$, where nodes represent clinical and operational factors such as patient inflow, bed occupancy, staff availability, and external environmental indicators. The edges evolve over time to reflect changing correlations. This design allows the model to effectively capture dynamic bottlenecks and resource constraints analogous to supply-demand fluctuations.

The node embedding at time t is computed using a graph attention mechanism combined with temporal encoding:

$$h_i^t = \sigma\left(\sum_{j \in N(i)} \alpha_{ij}^t \mathbf{h}_j^{t-1}\right)$$

where α_{ij}^t represents the attention coefficient learned dynamically based on temporal similarity and operational dependency. The temporal evolution is further modeled using a Transformer-based encoder:

$$Z_t = \text{Transformer}(H_{t-k:t})$$

The final fused representation is obtained as:

$$F_t = \lambda_1 Z_t + \lambda_2 G_t$$

where λ_1 and λ_2 are adaptive weighting parameters learned during training. This fusion allows MedAgentOps to jointly capture both structural relationships and long-term temporal dynamics in emergency operations.

3.3 Shared Medical Memory (SMM)

A key limitation of existing forecasting and decision systems is the lack of persistent memory across time. To address this, we introduce the Shared Medical Memory (SMM) module, which stores both short-term operational states and long-term historical experiences across all agents.

We define two memory components: short-term memory M_s and long-term memory M_l . The short-term memory stores recent system states:

$$M_s = \{F_{t-1}, F_{t-2}, \dots, F_{t-\omega}\}$$

while long-term memory encodes historical patterns:

$$M_l = \{H^{(1)}, H^{(2)}, \dots, H^{(N)}\}$$

Memory retrieval is performed using similarity-based attention:

$$\alpha_i = \frac{\exp(\text{sim}(q, M_i))}{\sum_j \exp(\text{sim}(q, M_j))}$$

and the retrieved context is computed as:

$$M_{\text{retrieved}} = \sum_i \alpha_i M_i$$

Inspired by recent advances in deep learning-based semantic recommendation database systems where knowledge representations are efficiently queried for personalized retrieval [13], this memory mechanism allows MedAgentOps to generalize across different operational scenarios, such as seasonal peaks, epidemic outbreaks, and emergency surges.

3.4 Agent Graph Reasoning Module (AGRM)

To enable structured collaboration among heterogeneous agents, we propose the Agent Graph Reasoning Module (AGRM), which models all agents as nodes in a reasoning graph $GA = (A, R)$, where A denotes agents and R denotes interaction relationships.

Each agent performs a specialized task, such as forecasting, scheduling, or resource allocation. The interaction between agents is modeled using graph attention:

$$\beta_{ij} = \frac{\exp(\phi(a_i, a_j))}{\sum_k \exp(\phi(a_i, a_k))}$$

where $\phi(\cdot)$ is a learnable compatibility function.

The updated agent state is:

$$a'_i = \sum_{j \in N(i)} \sigma(\beta_{ij} W a_j)$$

This mechanism allows dynamic adaptation of reasoning paths depending on the operational context. For instance, when patient inflow increases sharply, the prediction agent assigns higher importance to scheduling and resource agents, enabling rapid response coordination.

3.5 Memory-Augmented Multi-Agent Reasoning in MAGR-Net

The MAGR-Net integrates temporal representation, memory retrieval, and agent reasoning into a unified architecture. The overall inference process can be formulated as:

$$Y_t = G(F_t, M_{retrieved}, A_t)$$

where G denotes the global reasoning function, F_t represents fused temporal features, $M_{retrieved}$ denotes memory context, and A_t represents agent states.

The prediction output for emergency inflow is:

$$\hat{y}_{t+1:t+T} = f_{pred}(F_t, M_{retrieved})$$

while the resource optimization policy is defined as:

$$\pi^* = \arg \max_{\pi} E \left[\sum_t \gamma^t (R_t^{eff} - R_t^{cost} - R_t^{wait}) \right]$$

where the reward function balances efficiency, operational cost, and patient waiting time.

3.6 Multi-Objective Reinforcement Learning for Resource Optimization (MOPPO)

To optimize hospital operations, we adopt a Multi-Objective Proximal Policy Optimization (MOPPO) strategy. The state space includes ED occupancy, staff availability, and predicted inflow:

$$s_t = [\hat{y}_t, b_t, d_t, r_t]$$

where b_t , d_t , and r_t denote beds, doctors, and resources respectively.

The action space includes dynamic scheduling decisions:

$$a_t = \{\text{allocate beds, adjust staff, activate emergency resources}\}$$

The reward function is defined as:

$$R_t = \alpha \cdot (-W_t) + \beta \cdot U_t + \gamma \cdot (-C_t)$$

where W_t is waiting time, U_t is resource utilization, and C_t is operational cost.

Policy optimization follows:

$$\theta^* = \arg \max_{\theta} E_{\pi_{\theta}} \left[\sum_t R_t \right]$$

This enables adaptive decision-making under highly dynamic emergency conditions.

3.7 Distributed Agent Infrastructure and Execution Mechanism

MedAgentOps is implemented as a distributed agent system where each agent operates asynchronously over a shared message bus. Communication is event-driven, allowing real-time coordination across prediction, reasoning, and optimization modules. To reduce inference latency, we introduce a shared KV-cache mechanism and dynamic scheduling policy:

$$Score(a_i) = \alpha U_i + \beta Q_i + \gamma P_i$$

where U_i is urgency, Q_i is queue length, and P_i is predicted impact. Agents with higher scores are allocated GPU resources preferentially. Moreover, the decentralized nature of MedAgentOps makes it highly adaptable to edge computing paradigms. For perception agents that handle localized visual triage or real-time clinical image extraction, the system pipeline is designed to be compatible with low-power hardware accelerators (e.g., FPGAs) [14]. This multi-tiered execution mechanism—from hardware-accelerated feature extraction at the edge to graph-based reasoning on central GPUs—ensures robust scalability. This infrastructure ensures scalability and efficiency when deployed in large-scale hospital environments with multiple concurrent emergency streams.

4. Experiment

4.1 Dataset Preparation

In this study, we utilize a real-world emergency department (ED) operational dataset collected from a large tertiary-care hospital's information management system. The dataset is derived from the hospital's Electronic Health Records (EHR) system and emergency department admission logs, covering a continuous period of 3 years from January 2021 to December 2023. It records both patient-level arrival information and system-level operational states, enabling comprehensive modeling of ED flow dynamics and resource utilization patterns. The dataset contains 182,460 emergency visit records, aggregated at hourly resolution for forecasting tasks. Each record includes multi-dimensional features describing patient inflow, environmental conditions, and hospital operational status. Specifically, the dataset integrates temporal attributes (hour, day of week, month, holiday indicator), clinical demand indicators (number of arrivals, triage severity level distribution), and hospital resource variables (available beds, occupied beds, number of on-duty physicians and nurses). In addition, external contextual factors such as weather conditions and regional epidemic indicators are incorporated to enhance predictive capability.

A summary of the key dataset features is presented in Table 1.

Table 1. Overview of ED Dataset Features.

Feature Category	Feature Name	Description
Temporal Features	Hour, Day, Month	Time indicators for periodicity modeling
Demand Features	Patient Arrival Count	Number of incoming ED patients per hour
Clinical Features	Triage Level Distribution	Severity classification of patients
Resource Features	Bed Availability	Number of available emergency beds
Resource Features	Staff Allocation	Number of active doctors and nurses
External Factors	Weather Condition	Temperature, humidity, and weather state
External Factors	Epidemic Index	Regional infection intensity indicator

Overall, this dataset provides a high-resolution spatio-temporal view of emergency department operations, making it well-suited for evaluating the proposed MedAgentOps framework in both forecasting accuracy and resource optimization effectiveness.

4.2 Experimental Setup

The experimental evaluation is conducted on a real-world emergency department dataset collected from a large tertiary hospital, spanning three years of continuous operational records. The dataset is aggregated at an hourly granularity and includes patient arrival counts, resource utilization states, and external environmental factors. To ensure a fair comparison, the dataset is split into training,

validation, and testing sets with a ratio of 7:1:2, following standard time-series forecasting protocols without data leakage. All models are implemented using PyTorch 2.0 and trained on a server equipped with NVIDIA A100 GPUs (40GB memory), Intel Xeon CPUs, and 256GB RAM. The optimization uses the Adam optimizer with an initial learning rate of $1e-4$ and cosine decay scheduling. For MedAgentOps, we deploy a distributed multi-agent environment where each agent runs asynchronously and communicates via a shared message bus. The maximum sequence length is set to 96 hours, and the batch size is fixed at 64. Early stopping is applied based on validation loss with a patience of 10 epochs to avoid overfitting.

4.3 Evaluation Metrics

To comprehensively evaluate forecasting accuracy and operational efficiency, we adopt multiple widely used metrics in time-series prediction and healthcare optimization tasks. Prediction performance is measured using Mean Absolute Error (MAE), Mean Absolute Percentage Error (MAPE), and Root Mean Square Error (RMSE), which evaluate absolute deviation, relative error, and variance-sensitive deviation respectively. The forecasting task is defined as predicting future emergency department patient inflow over a multi-step horizon. In addition, to evaluate resource optimization effectiveness, we introduce operational metrics including average patient waiting time (AWT), resource utilization rate (RUR), and cost efficiency (CE). These metrics reflect the ability of the system to translate forecasting results into effective scheduling decisions. Higher RUR and lower AWT indicate better system performance, while CE evaluates the trade-off between operational cost and service quality.

4.4 Results

The results in Table 2 demonstrate that the proposed MAGR-Net significantly outperforms all baseline models in emergency department flow forecasting. Specifically, MAGR-Net achieves an MAE of 10.21, reducing prediction error by 26.6% compared to GCN-LSTM (13.92) and by 44.5% compared to LSTM (18.42). In terms of RMSE, MAGR-Net achieves the lowest value of 14.33, indicating improved robustness against extreme fluctuations in patient arrival patterns. Furthermore, the MAPE of 8.73% demonstrates superior relative forecasting accuracy compared to TFT (11.37%) and Transformer (12.85%). These improvements are primarily attributed to the integration of Dynamic Temporal Graph Fusion and Shared Medical Memory, which enable richer temporal dependency modeling and historical knowledge utilization. Overall, the results confirm that memory-augmented graph reasoning significantly enhances predictive performance in complex healthcare environments.

Table 2. Forecasting Performance Comparison.

Model	MAE ↓	RMSE ↓	MAPE (%) ↓
LSTM	18.42	24.87	14.96
Transformer	16.31	22.14	12.85
TFT	14.76	20.08	11.37
GCN-LSTM	13.92	19.41	10.54
MAGR-Net (Ours)	10.21	14.33	8.73

Table 3 presents the performance of different models in emergency department resource optimization. The proposed MedAgentOps framework achieves the best overall performance across all evaluation metrics. Specifically, it reduces average patient waiting time to 32.5 minutes, which is 18.7% lower than GNN-RL (39.6 minutes) and 47.9% lower than rule-based scheduling (62.4 minutes). In terms of resource utilization, MedAgentOps achieves 91.3%, significantly outperforming MARL (82.1%) and PPO-based methods (78.5%). Moreover, cost efficiency reaches 0.91, indicating a superior balance between operational cost and service quality. These improvements stem from the integration of Multi-Objective Proximal Policy Optimization and Agent Graph Reasoning Module, which enable adaptive coordination among multiple specialized agents. The results strongly demonstrate that MedAgentOps is highly effective in translating predictive intelligence into actionable operational decisions in dynamic emergency environments.

Table 3. Resource Optimization Performance Comparison.

Model	Waiting Time ↓	Utilization ↑	Cost Efficiency ↑
Rule-based Scheduling	62.4 min	71.2%	0.68
RL (PPO)	48.7 min	78.5%	0.74
MARL	42.3 min	82.1%	0.79
GNN-RL	39.6 min	84.7%	0.83
MedAgentOps (Ours)	32.5 min	91.3%	0.91




Figure 2. Agent Interaction Graph Evolution in MedAgentOps.

Figure 2 illustrates the evolution of interaction importance weights among the four core agents in the proposed MedAgentOps framework during training. As training progresses from Epoch 1 to Epoch 20, all agents exhibit an overall upward trend with moderate fluctuations, indicating that the collaborative reasoning mechanism gradually learns more effective inter-agent communication patterns. The Prediction Agent consistently maintains the highest interaction weight, increasing from 0.42 at Epoch 1 to 0.81 at Epoch 20. This observation is consistent with the forecasting results

reported in Table 2, where MAGR-Net achieves the best predictive performance with a MAPE of 8.73%, MAE of 10.21, and RMSE of 14.33. The Scheduling Agent rises from 0.35 to 0.77, reflecting its growing contribution to resource allocation decisions. Similarly, the Resource Agent improves from 0.30 to 0.73, while the Knowledge Agent increases from 0.28 to 0.68, demonstrating the effectiveness of shared memory and knowledge retrieval mechanisms. The gradual strengthening of agent interactions corresponds to the resource optimization improvements shown in Table 3, where MedAgentOps reduces average waiting time to 32.5 minutes and increases resource utilization to 91.3%. These results confirm that enhanced multi-agent collaboration directly contributes to both forecasting accuracy and operational efficiency.

4.5 Discussion

The experimental results clearly indicate that integrating multi-agent collaborative reasoning with memory-augmented graph learning leads to substantial improvements in both forecasting accuracy and operational efficiency. Unlike traditional models that treat prediction and scheduling as separate tasks, MedAgentOps establishes a unified reasoning pipeline where prediction outputs directly influence downstream decision-making processes. Similar to how role-specialized agents enhance complex visual workflows through structured collaboration [15], this end-to-end coupling significantly reduces information loss and improves system responsiveness. Moreover, the proposed distributed architecture allows for the seamless integration of specialized diagnostic modules, such as deep learning models for rapid X-ray screening [16] and multi-scale attention-enhanced Transformer architectures for accurate vessel segmentation [17], and real-time CNN-Attention frameworks for rapid pathogen identification and virulence prediction [18], directly into the ED workflow to further reduce diagnostic delays. Furthermore, the inclusion of Shared Medical Memory enables the system to generalize across different operational scenarios, such as seasonal peaks and emergency outbreaks. The Agent Graph Reasoning Module further enhances adaptability by dynamically adjusting inter-agent dependencies based on contextual changes. This allows the system to shift reasoning priorities in real time, improving robustness under highly volatile conditions. Overall, the superior performance of MedAgentOps demonstrates the effectiveness of combining temporal modeling, graph reasoning, reinforcement learning, and distributed multi-agent infrastructure within a single unified framework for intelligent healthcare operation optimization.

5. Conclusions

This study investigated the challenging problem of emergency department (ED) flow forecasting and resource optimization in modern healthcare systems by proposing MedAgentOps, a distributed multi-agent collaborative reasoning infrastructure built upon the Memory-Augmented Graph Reasoning Network (MAGR-Net). Unlike conventional approaches that treat patient flow prediction and resource scheduling as independent tasks, the proposed framework integrates forecasting, memory-enhanced knowledge management, graph-based agent collaboration, and reinforcement learning-driven optimization into a unified end-to-end decision-making architecture. Through the Dynamic Temporal Graph Fusion (DTGF) module, Shared Medical Memory (SMM), Agent Graph Reasoning Module (AGRM), and Multi-Objective Proximal Policy Optimization (MOPPO) strategy, MedAgentOps enables intelligent coordination among multiple specialized agents to support real-time hospital operational management. Comprehensive experiments were conducted on a real-world emergency department dataset containing multi-source clinical, operational, and environmental information. The experimental results demonstrate that the proposed framework achieves superior performance compared with several state-of-the-art baseline methods, including LSTM, Transformer, Temporal Fusion Transformer (TFT), GCN-

LSTM, PPO, MARL, and GNN-RL. Specifically, MAGR-Net achieved the best forecasting performance with an MAE of 10.21, an RMSE of 14.33, and a MAPE of 8.73%, significantly outperforming all competing models. In addition, the MedAgentOps framework reduced the average patient waiting time to 32.5 minutes, representing an approximately 17.9% improvement over the strongest baseline model, while increasing resource utilization to 91.3% and achieving a cost-efficiency score of 0.91. These results indicate that integrating memory-augmented graph learning with multi-agent collaborative reasoning can effectively improve both forecasting accuracy and operational decision quality in complex healthcare environments. From an application perspective, the proposed framework provides a scalable and intelligent solution for hospital emergency management, healthcare resource planning, and smart medical operation systems. By enabling adaptive coordination among heterogeneous agents and leveraging historical operational knowledge, MedAgentOps offers strong potential for deployment in large-scale hospitals facing dynamic patient demand and resource constraints. Despite the encouraging results, several limitations remain. First, the current study focuses primarily on a single-hospital dataset, which may limit the generalizability of the model across different healthcare systems and geographical regions. Second, the agent collaboration mechanism is evaluated under relatively stable operational settings and may require further validation under large-scale emergency events or pandemic scenarios. Future research will explore cross-hospital collaborative learning, federated multi-agent architectures, large language model-enhanced medical agents, and multimodal data integration involving clinical notes, medical imaging, and public health information. Furthermore, incorporating causal reasoning and adaptive knowledge evolution mechanisms may further improve the robustness and interpretability of intelligent healthcare decision-making systems.

References

- Graves A. Long short-term memory[J]. Supervised sequence labelling with recurrent neural networks, 2012: 37-45.
- Vaswani A, Shazeer N, Parmar N, et al. Attention is all you need[J]. Advances in neural information processing systems, 2017, 30.
- Lim B, Arik S Ö, Loeff N, et al. Temporal fusion transformers for interpretable multi-horizon time series forecasting[J]. International journal of forecasting, 2021, 37(4): 1748-1764.
- Morid M A, Sheng O R L, Dunbar J. Time series prediction using deep learning methods in healthcare[J]. ACM Transactions on Management Information Systems, 2023, 14(1): 1-29.
- Mnih V, Kavukcuoglu K, Silver D, et al. Human-level control through deep reinforcement learning[J]. nature, 2015, 518(7540): 529-533.
- Schulman J, Wolski F, Dhariwal P, et al. Proximal policy optimization algorithms[J]. arXiv preprint arXiv:1707.06347, 2017.
- Li Y, Mao C, Huang K, et al. Deep reinforcement learning for efficient and fair allocation of health care resources[J]. arXiv preprint arXiv:2309.08560, 2023.
- Kipf T N, Welling M. Semi-supervised classification with graph convolutional networks[J]. arXiv preprint arXiv:1609.02907, 2016.
- Veličković P, Cucurull G, Casanova A, et al. Graph attention networks[J]. arXiv preprint arXiv:1710.10903, 2017.
- Rossi E, Chamberlain B, Frasca F, et al. Temporal graph networks for deep learning on dynamic graphs[J]. arXiv preprint arXiv:2006.10637, 2020.
- Lowe R, Wu Y I, Tamar A, et al. Multi-agent actor-critic for mixed cooperative-competitive environments[J]. Advances in neural information processing systems, 2017, 30.
- Hu Z, Hu M, Li Z. DynSupplyNet: A Dynamic Graph Neural Network with Temporal Fusion for Supply Chain Risk Prediction and Propagation Analysis[J]. Journal of Computer Science and Frontier Technologies, 2026, 3(2): 1-13.
- Wang T, Zhang K, Liu K. A Knowledge Graph and Deep Learning-Based Semantic Recommendation Database System for Advertisement Retrieval and Personalization[J]. arXiv preprint arXiv:2601.00833, 2025.
- Wang Y. Low-power design of advanced image processing algorithms under fpga in real-time applications[C]//2024 IEEE 4th International Conference on Power, Electronics and Computer Applications (ICPECA). IEEE, 2024: 1080-1084.
- Zhang B, Wang X, Kao H, et al. Consensus Matrix: A Role-Specialized Multi-Agent Framework for Structured Collaborative Decision-Making in Agentic Visual Media Workflows[C]//Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2026: 4613-4622.
- Yao Y, Li Y, Zheng S, et al. Deep Learning Models for Multi-class Pneumonia Detection in Chest X-rays: A Comparative Study of VGG16, MobileNet, and ResNet152[J]. Computer Simulation in Application, 2025, 3(1): 66-73.
- Yao Y, Li Y, Zheng S, et al. MSA-TransUNet: A Multi-Scale Attention Enhanced Transformer-UNet Architecture for Accurate Vessel Segmentation and Visualization in Medical Imaging[J]. Journal of Medicine and Life Sciences, 2025, 1(4): 48-57.
- Xu H, Ge J, Ou Y. DentaScope-AI: A Real-Time k-mer-Driven CNN-Attention Framework for Rapid Pathogen Identification and Virulence Prediction in Endodontic Infections Using Nanopore Sequencing[J]. Frontiers in Healthcare Technology, 2026, 3(1): 112-122.