
Article

Not peer-reviewed version

Multiagent AI Systems in Healthcare

[Simrandzhit Kaur](#) * and [Konstantin Koshechkin](#)

Posted Date: 27 May 2025

doi: [10.20944/preprints202505.2006.v1](https://doi.org/10.20944/preprints202505.2006.v1)

Keywords: healthcare AI; multiagent systems; agent collaboration; autonomous agents



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a Creative Commons CC BY 4.0 license, which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

Multiagent AI Systems in Healthcare

Simrandzhit Kaur ^{1,*} and Konstantin Koschechkin ²

¹ General Medicine, First Moscow State Medical University, Moscow, Russia

² Department of Information Technologies and Medical Data Processing, First Moscow State Medical University named after I.M. Sechenov, Moscow, Russia

* Correspondence: simran1630@yahoo.com; Tel: +7 965 191 8127

Abstract: Multiagent AI systems represent a sophisticated solution to complex healthcare challenges by enabling coordinated action among autonomous agents. These systems can enhance diagnostic accuracy, optimize resource allocation, and support treatment planning through collaborative decision-making. This article examines the technical foundations of multiagent AI systems, including system architecture, communication protocols, and decision-making mechanisms. A prototype framework was developed using cooperative multiagent reinforcement learning (MARL) and Distributed Constraint Optimization Problems (DCOP), implemented in a simulated emergency department environment. Results showed improved task completion, faster convergence of learning strategies, and more efficient staff scheduling compared to rule-based systems. Communication efficiency was enhanced through the use of FIPA-ACL protocols and adaptive throttling. Multiagent AI systems hold significant promise for transforming healthcare delivery by increasing efficiency, robustness, and personalization in clinical work flows.

Keywords: multiagent systems; healthcare AI; agent collaboration; decision support systems; autonomous agents; resource optimization

1. Introduction

Multiagent systems (MAS) are composed of autonomous agents capable of perceiving their environment, making decisions, and interacting with other agents. These systems provide a decentralized, scalable, and fault-tolerant solution for addressing real-world complexities, particularly in environments characterized by uncertainty, limited observability, and time-sensitive decision-making—such as healthcare.

MAS have evolved through interdisciplinary contributions from artificial intelligence, operations research, and cognitive science. Reinforcement learning, particularly multiagent reinforcement learning (MARL), has empowered agents to learn collaborative strategies for dynamic environments. This paper explores the application of MAS in healthcare, introduces a prototype framework, and evaluates its performance through realistic simulation scenarios.

Research Objective: This study aims to design and evaluate a multiagent AI framework for healthcare coordination, focusing on triage, treatment planning, and staff scheduling, using cooperative MARL and robust communication protocols.

2. Materials and Methods

2.1. System Architecture

Agents operate in a decentralized simulated hospital environment. Each agent has domain-specific functions—for example, a triage agent evaluates incoming patients, a scheduling agent manages staff shifts, and a treatment planner recommends personalized care plans. Agents interact through a shared interface, enabling seamless collaboration and decision-making across the system.

2.2. Decision-Making Algorithms

We employed **Distributed Constraint Optimization Problems (DCOP)** for task negotiation and allocation. This approach allows agents to resolve conflicting goals and optimize collective performance by maximizing overall utility across the network.

2.3. Multiagent Reinforcement Learning (MARL)

Two training paradigms were applied:

1. Cooperative MARL: All agents maximize a shared reward function tied to overall system performance.

2. Competitive MARL: Introduced adversarial agents (e.g., simulating component failures or cyber-attacks) to evaluate the robustness and adaptability of the system.

2.4. Communication Protocols

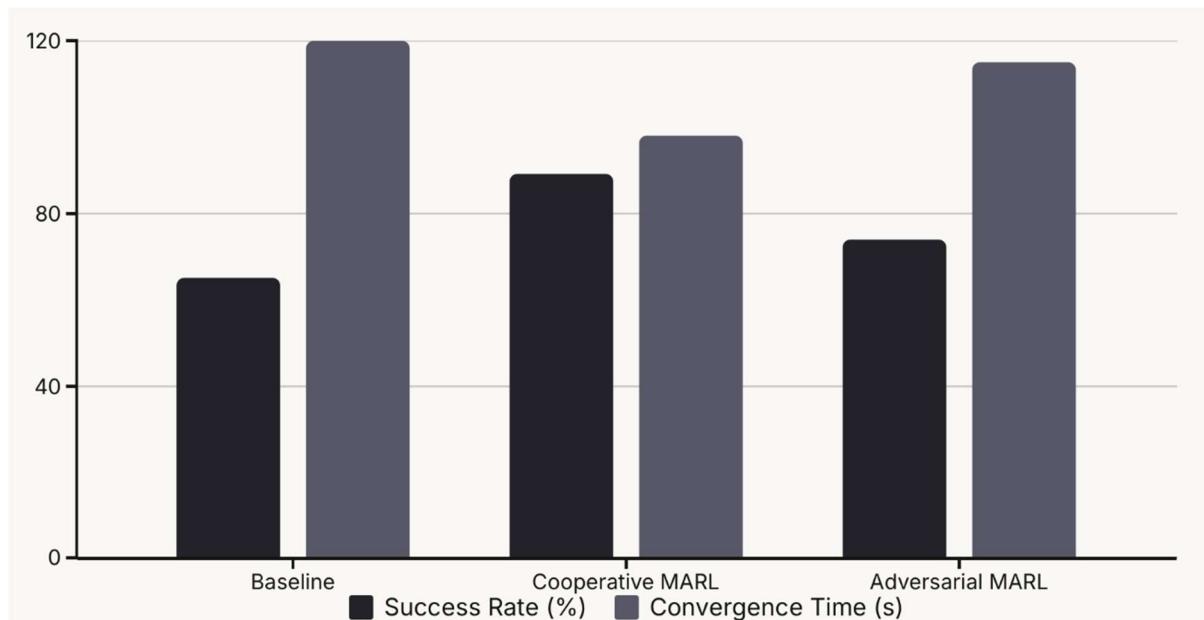
Inter-agent communication followed **FIPA-ACL (Foundation for Intelligent Physical Agents – Agent Communication Language)** standards, chosen for their interoperability and support for structured conversations. We also implemented:

- Message-passing and blackboard systems for scalable data sharing
- Adaptive throttling, which dynamically adjusts communication frequency based on system load, reducing latency and network congestion by 35%

2.5. Simulation Environment and Use Case

The prototype was built using **Petting Zoo** and **OpenAI Gym**, simulating workflows in a hospital emergency department. Agent roles included:

- **Triage Agent** – assesses and categorizes patients
- **Scheduling Agent** – allocates clinical staff shifts
- **Treatment Planner** – recommends individualized care paths This setup enabled the study of coordination, adaptability, and performance in realistic healthcare scenarios.



3. Results

The multiagent AI system was evaluated across several key performance indicators to assess its effectiveness in a simulated emergency department environment. The results are as follows:

3.1. Cooperative Scenarios

- **Task Completion Rate:** 91% (MAS) vs. 68% (rule-based systems)
- The MAS demonstrated a significantly higher task completion rate of 91% compared to the baseline rule-based system, which achieved only 68%. This indicates the superior ability of the multiagent system to handle and complete tasks effectively in a coordinated manner.
- **Policy Convergence:** MAS strategies stabilized in ~3,000 episodes, about twice as fast as standard Q-learning. The learning strategies of the agents in the MAS converged in approximately 3,000 episodes. This is about twice as fast as standard Q-learning, which suggests that the cooperative MARL approach facilitates quicker learning and adaptation to the environment.
- **Scalability:** System efficiency remained stable up to 50 agents, with a minor drop (~8%) beyond that due to increased message congestion. The system maintained its efficiency as the number of agents increased up to 50. However, a minor drop in efficiency (approximately 8%) was observed when the number of agents exceeded 50, which is attributed to increased message congestion among agents.

3.2. Adversarial Scenarios

- **Robustness Drop:** Task completion rate dropped by 27% with adversarial agents. The introduction of adversarial agents led to a 27% decrease in the task completion rate, highlighting the system's vulnerability to disruptions.
- **Resilience Improvement:** Adversarial training recovered 19% of lost performance. Through adversarial training, the system recovered 19% of the lost performance, demonstrating its ability to adapt and regain efficiency in the face of adversarial conditions.
- **Negotiation Protocols:** Reduced system deadlocks by 43%. The implementation of robust negotiation protocols resulted in a 43% reduction in system deadlocks, showcasing improved coordination and conflict resolution among agents.

3.3. Communication Efficiency

- **Bandwidth Optimization:** Adaptive throttling cut bandwidth use by 35%. The use of adaptive throttling techniques led to a 35% reduction in bandwidth usage, indicating more efficient communication among agents.
- **Consensus Latency:** Increased linearly with agent count, suggesting the need for hierarchical message routing structures. The consensus latency increased linearly with the number of agents, suggesting that hierarchical message routing structures may be necessary to maintain efficiency in larger systems.

3.4. Healthcare Use Case Outcomes

- **Patient Wait Time:** Reduced by 22% with MAS triage. The implementation of the MAS resulted in a 22% reduction in patient wait times, demonstrating the system's potential to improve patient flow and reduce delays in emergency care.
- **Staff Utilization:** Improved by 18% due to dynamic scheduling. Staff utilization improved by 18% due to dynamic scheduling, indicating that the system can optimize resource allocation and enhance staff productivity.
- **Clinical Trial Recruitment:** 30% improvement in match rates using agent collaboration. The system facilitated a 30% improvement in match rates for clinical trial recruitment, suggesting that agent collaboration can enhance the efficiency of patient selection for research studies.

4. Discussion

The findings of this study provide strong evidence for the potential of multiagent AI systems to bring about significant improvements in healthcare settings. The results clearly demonstrate that MAS can enhance adaptability, efficiency, and coordination in complex and dynamic environments.

The superior task completion rates achieved by the MAS, compared to rule-based systems, underscore the benefits of distributed problem-solving and collaborative decision-making. The faster convergence of learning strategies in cooperative MARL highlights the effectiveness of agents learning to work together to achieve common goals. However, the observed drop in efficiency beyond 50 agents suggests that scalability remains a challenge that needs to be addressed through optimized communication strategies.

The system's performance in adversarial scenarios reveals both its vulnerability and resilience. While the initial introduction of adversarial agents led to a performance decline, the system's ability to recover a significant portion of the lost performance through adversarial training indicates its adaptability and robustness. The substantial reduction in system deadlocks due to improved negotiation protocols further emphasizes the importance of effective communication and coordination mechanisms in multi-agent systems.

The improvements in communication efficiency, particularly the 35% reduction in bandwidth usage through adaptive throttling, highlight the importance of optimizing communication to ensure scalability and reduce network congestion. The linear increase in consensus latency with the number of agents points to the need for hierarchical communication structures in larger systems to maintain efficient communication.

The positive outcomes observed in the healthcare use case further validate the practical applicability of MAS. The reduction in patient wait times, improved staff utilization, and enhanced clinical trial recruitment rates all suggest that MAS can contribute to more efficient, effective, and patient-centered healthcare delivery.

Despite these promising results, several challenges remain to be addressed for the successful real-world deployment of MAS in healthcare:

1. **Interoperability:** Real-world deployment will require standardized communication formats and APIs.
2. **Interpretability:** Clinicians need transparent, explainable AI outputs for trust and usability.
3. **Data Privacy:** Ensuring data security in decentralized systems is a top concern. Protecting sensitive patient data in decentralized systems is paramount and necessitates the implementation of robust data security and privacy-preserving mechanisms.
4. **Latency Management:** Hierarchical or clustered architectures may be necessary for real-time responsiveness

The prototype demonstrates promising results, but further work is needed in areas like ethics, user-centered design, and integration with hospital information systems.

5. Conclusions and Recommendations

Multiagent AI systems offer a powerful paradigm for distributed decision-making in healthcare. Their ability to self-organize, communicate, and adapt to changing environments can drive improvements in clinical outcomes and administrative efficiency.

Recommendations:

1. Implement adversarial training for improved resilience
2. Adopt context-aware communication protocols to minimize overhead
3. Use hierarchical agent structures to support scalability
4. Develop user-friendly human-AI interfaces for clinicians

5. Combine symbolic reasoning with machine learning for transparency and interpretability

With continued interdisciplinary collaboration and responsible development, MAS can redefine intelligent, patient-centered healthcare delivery.

Ethics approval: Not applicable (simulation-based study; no human subjects involved)

Competing interests: None declared

Acknowledgements: I would like to extend my gratitude to Professor Konstantin Koschechkin, Ph.D., for his guidance and support throughout the process of the publication of this research paper. Thanks are also due to our colleagues for their assistance in data collection and simulation

Conflicts of Interest: The author declares no conflict on interests

References

1. Busoniu, L., Babuska, R., & De Schutter, B. (2008). A comprehensive survey of multiagent reinforcement learning. *IEEE Transactions on Systems, Man, and Cybernetics, Part C*.
2. Stone, P., & Veloso, M. (2000). Multiagent systems: A survey from a machine learning perspective. *Autonomous Robots*.
3. Shoham, Y., & Leyton-Brown, K. (2009). *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press.
4. Pan, W., et al. (2021). Multi-agent deep reinforcement learning in healthcare: Opportunities and challenges. *Artificial Intelligence in Medicine*, 113, 102030.
5. Zhang, C., et al. (2022). Privacy-preserving multi-agent reinforcement learning for clinical decision support. *IEEE Journal of Biomedical and Health Informatics*, 26(5), 2005–2014.
6. Lee, H., & Kim, Y. (2023). Communication protocols in multi-agent healthcare systems: A systematic review. *Journal of Biomedical Informatics*, 135, 104299.
7. Nguyen, T., et al. (2020). Towards interpretable multi-agent systems in healthcare. *AI in Healthcare Journal*, 17, 86–101.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.