

“Coordinated AI agents for advancing healthcare”: A Review

Moritz, M., Topol, E., & Rajpurkar, P. (2025). Coordinated AI agents for advancing healthcare. *Nature Biomedical Engineering*, 9(5), 432–438. <https://doi.org/10.1038/s41551-025-01363-2>

Joseph Katz - 07/02/2025

Comment

<https://doi.org/10.1038/s41551-025-01363-2>

Coordinated AI agents for advancing healthcare

Michael Moritz, Eric Topol & Pranav Rajpurkar

 Check for updates

Decentralized yet coordinated networks of specialized artificial intelligence agents, multi-agent systems for healthcare (MASH), that excel in performing tasks in an assistive or autonomous manner within specific clinical and operational domains are likely to become the next paradigm in medical artificial intelligence.

AI agents function within the autonomous–assistive spectrum of tasks. Autonomous AI agents can work independently to collect and process data and take actions within their designated areas of expertise. For example, an autonomous AI agent in radiology might analyse medical images to detect signs of disease and automatically generate a report. Assistive AI agents support human decision-making by providing insights, recommendations or suggestions to healthcare providers, who then make the final decisions on the basis of their expertise and the AI-generated information. For example, a conversational agent powered by an LLM or a vision–language model trained on a massive corpus of biomedical literature and current guidelines can provide

What is an AI Agent?

AI Agent: A software program built on a large language model (like ChatGPT) that can understand and respond to human language, make decisions, search the internet or databases, write code, and interact with digital interfaces. It can either assist a human user or complete tasks on its own.


- **Assistive AI Agent:** Supports human users by offering recommendations, retrieving information, or helping interpret data. The human retains final decision-making.
 - Example: An AI that summarizes clinical guidelines for a physician or suggests lab follow-ups.
- **Autonomous AI Agent:** Capable of completing defined tasks entirely on their own, without needing human input at every step.
 - Example: An AI that drafts radiology reports or processes insurance claims automatically.

What is MASH (Multi-Agent System for Healthcare)?


Multi-agent systems for healthcare (MASH): A group of AI agents that each specialize in different healthcare tasks—like diagnosis, scheduling, or insurance—and work together as a team. Even though each agent runs independently (decentralization), they coordinate with one another to help clinicians and patients, either by assisting decisions or carrying out tasks on their own.



Multiagent model example illustration




MASH Example Illustration – A Patient's Journey



Why Use MASH?

- Healthcare is complex: tasks include diagnosis, scheduling, billing, treatment planning, etc.
- No single model can do all tasks equally well – use specialized agents
- Coordination allows for more efficient and personalized care
- Reflects real clinical cooperation



of Agents v Accuracy (Grade School Math 8k benchmark)¹

1. Du, Y., Zhu, Z., Weng, W., Zhou, Y., Mao, X., Pan, Z., ... & Liu, J. (2024). *Do we still need monolithic LLMs? An evaluation of modular LLMs on unseen tasks*. OpenReview.

<https://openreview.net/pdf?id=bgzUSZ8aeg>

Real-World Uses and Clinical v. Operational Agents

Clinical Agents

Supports medical decisions

- Diagnostic Assistants
- Treatment Suggestions
- Monitoring
- Medical Question Support

Operational Agents

Manages administrative and logistical tasks

- Scheduling Tools
- Bed Flow Management
- Insurance Processing
- Inventory/Staffing Tools



MASH Benefits

Author's Arguments

- Improved Safety/Privacy (Decentralization)
- Reliability with Cross-Checking Agents
- Efficiency and Cost/Resource Optimization
- Interpretability – Human-Readable
- Human-centered Design, Empathy

Additional Arguments

- Llama, Gemma, LM Studio Local downloads, privacy and security with hospital-patient information training sets
- Applications in Research: Donor Project
- Healthcare training: patient-doctor scenarios in medical education (MEDCO)²
- Health Equity/Accessibility: distribution of care
- Hyper-personalization: 24/7 individual monitoring, personal recommendations and plan creation³

MASH Limitations

Author's Arguments

- Coordination complexity: so many connections + error prone
- Cascade of Errors
- Infrastructure Burden and Costs
- Trust Concerns

Additional Arguments

- Amplified Discrimination: algorithmic bias and training data⁴
- Cloud APIs/Web Search Leaks (Princeton and HIPAA)⁵
- Loss of trust in human care/workforce impact
- Complete workflow disruption/rewriting education
- Accountability, who is to blame?
- Alert Fatigue: too many alarms sounded

4. Pavuluri, S., Sangal, R., Sather, J., & Taylor, R. A. (2024). Balancing act: the complex role of artificial intelligence in addressing burnout and healthcare workforce dynamics. *BMJ health & care informatics*, 31(1), e101120. <https://doi.org/10.1136/bmjhci-2024-101120>

5. Price School of Public Policy. (2023, October 23). *Why doctors using ChatGPT are unknowingly violating HIPAA*. University of Southern California. <https://priceschool.usc.edu/news/why-doctors-using-chatgpt-are-unknowingly-violating-hipaa/>

Overall Conclusions/Purpose (Benefit v. Harm)

Author

- MASH is a **promising next step** in clinical AI: scalable, modular, and task-specialized
- Enables **safe, personalized, efficient care** through teamwork-like agent collaboration
- Offers **greater reliability** than generalist LLMs by avoiding single-model failure
- Designed for flexibility: **agents can be added, swapped, or specialized** as needs evolve
- Not a plug-and-play solution: requires **human oversight, coordination protocols**, and clear regulation
- Could democratize healthcare access, especially for **under-resourced settings**
- Needs **interdisciplinary development** (clinicians, engineers, ethicists, policymakers)

Additions

- While MASH offers scalability, it may **require more institutional effort** than authors suggest—most hospitals aren't ready for this level of system integration.
- Authors focus on technical promise, but **trust, liability, and workflow culture** may be bigger barriers than model performance.
- The framework is promising, but **clinician training and AI literacy** are under-emphasized—without these, safe deployment is unlikely.
- The paper envisions equity, but **real-world rollout may deepen digital divides** if underserved hospitals can't afford or support MASH infrastructure.

Citations

1. Du, Y., Zhu, Z., Weng, W., Zhou, Y., Mao, X., Pan, Z., ... & Liu, J. (2024). *Do we still need monolithic LLMs? An evaluation of modular LLMs on unseen tasks*. OpenReview. <https://openreview.net/pdf?id=bgzUSZ8aeg>
2. Wei, H., Qiu, J., Yu, H., & Yuan, W. (2024). *MEDCO: Medical education copilots based on a multi-agent framework*. arXiv. <https://doi.org/10.48550/arXiv.2408.12496>
3. Pariveda Solutions. (2024, April 9). *Multi-agent AI: A game changer for healthcare innovation*. Pariveda. <https://parivedasolutions.com/perspectives/multi-agent-ai-a-game-changer-for-healthcare-innovation/>
4. Pavuluri, S., Sangal, R., Sather, J., & Taylor, R. A. (2024). Balancing act: the complex role of artificial intelligence in addressing burnout and healthcare workforce dynamics. *BMJ health & care informatics*, 31(1), e101120. <https://doi.org/10.1136/bmjhci-2024-101120>
5. Price School of Public Policy. (2023, October 23). *Why doctors using ChatGPT are unknowingly violating HIPAA*. University of Southern California. <https://priceschool.usc.edu/news/why-doctors-using-chatgpt-are-unknowingly-violating-hipaa/>