

Multi-Agent AI Systems in Healthcare: Systematic Evidence Synthesis Via PRISMA of Clinical Decision Support Systems, Robotic Interventions, and Critical Care

Oto Peter Ode¹, Ndukwe Chika Kalu², Saleh Abbas², Arslan Arshad², Adagboyi Isaac Inalegwu², Konstantin Koshechkin³

¹Department of Public Health, I.M. Sechenov First Moscow State Medical University

²Department of Epidemiology and Evidence-Based Medicine, I.M. Sechenov First Moscow State Medical University

³Department of Information and Internet Technologies, Digital Health Institute. I.M. Sechenov First Moscow State Medical University

DOI: <https://doi.org/10.51583/IJLTEMAS.2025.140500080>

Received: 01 June 2025; Accepted: 05 June 2025; Published: 19 June 2025

Abstract: Multi-agent artificial intelligence systems (MAS) is transforming the healthcare sector by enabling intelligent decision-making, robotic interventions, and critical care management. Despite growing interest, a comprehensive synthesis of their real-world applications, performance, and limitations is lacking. This systematic review aims to evaluate the deployment of MAS in healthcare across three key domains: Clinical Decision Support Systems (CDSS), Robotic Interventions, and Critical Care Monitoring. It also addresses ethical and legal considerations, offers a comparative framework, and identifies research gaps. A systematic search of PubMed, IEEE Xplore, Scopus, and Web of Science was conducted following PRISMA guidelines. A total of 150 records were retrieved. After duplicate removal and screening, 32 studies met inclusion criteria. The Joanna Briggs Institute (JBI) Critical Appraisal Checklist guided quality assessment. Data were extracted using a structured form and organized by topic. Among the 32 studies included, 12 focused on MAS-based CDSS, 11 on robotic interventions, and 9 on MAS in critical care. Most MAS demonstrated improved diagnostic accuracy, aided real-time decision support, and enhanced patient monitoring. However, over 60% involved practical models and lacked clinical validation. Ethical concerns, particularly around autonomy, data privacy, transparency, and legal accountability were sometimes overlooked. Only 7 studies discussed ethical or legal implications in depth. A comparative framework was developed, comparing MAS tools based on scalability, clinical validation, autonomy, human-AI interaction, and data transparency. No doubt, MAS hold transformative potential across healthcare domains. However, real-world deployment remains unclear. Greater emphasis on clinical validation, ethical design, and regulatory frameworks is needed. This review provides suggestions for future research and responsible adoption of MAS in healthcare.

Keywords: Multi-agent systems, artificial intelligence, healthcare robotics, critical care, clinical decision support systems, ethics, legal framework, PRISMA, cc.

I. Introduction

More than just a trend, the impact of Artificial Intelligence (AI) is becoming widespread. From transportation to education and healthcare, it is clear that AI is here to stay. In the health sector, AI is changing how services are delivered and is helping to improve the quality of care. Its use can already be seen in many areas of healthcare. However, the growing use of AI through Multi Agent Systems (MAS) in Clinical Decision Support Systems (CDSS), robotic procedures, and critical care monitoring deserves close attention.

MAS involves several AI agents working together to perform tasks for a user or another system. Each agent has its own features, but they work as a team to reach common goals and produce the desired results [1]. These systems are useful for handling large and complex tasks that may involve hundreds or even thousands of agents [2].

MAS provides many benefits over regular AI models. By sharing information and working toward shared goals, MAS allows for better decision-making, improved scale, and smarter responses based on the situation [3]. They are used in many areas, from sorting and diagnosing patients in emergencies to helping robotic surgery teams and monitoring patients in intensive care. But because they are used in so many ways, it can be hard to measure how well they work, how reliable they are, and whether they follow ethical rules.

Even though MAS can improve accuracy and save time in healthcare, there are still concerns about how well they perform in real-life situations, especially in high-risk settings [4]. Many of these systems are still being tested and are not yet fully used in everyday medical practice. Also, important ethical and legal issues like patient rights, data protection, responsibility, and informed consent are often not well handled in research [5]. This brings up serious concerns about trust, openness, and patient safety when MAS are used in direct care.

Although some studies have looked at MAS in specific areas, there is still no full and clear review of how they are built, how they work, and what their limits are. This review aims to fill that gap by looking at how MAS are used in three important areas: CDSS,

robotic procedures, and monitoring in critical care. It also looks at how ethical and legal issues are treated and offers a simple way to compare MAS tools, to guide future use in a safe and responsible way.

II. Method

Review Protocol

This systematic review was conducted in alignment with the PRISMA (Preferred Reporting Items for Systematic Reviews and Meta-Analyses) guidelines to ensure methodological rigor, transparency, and reproducibility in the review process. The protocol was designed to clearly define the review objectives, inclusion criteria, search strategy, and quality assessment approach before the data collection began.

Eligibility Criteria

We included studies that met the following criteria:

- Focused on multi-agent systems (MAS) used in healthcare
- Addressed at least one of these areas: Clinical Decision Support Systems (CDSS), robotic interventions, or critical care
- Were published in peer-reviewed journals between 2014 and 2024
- Provided evaluation or implementation results (quantitative, qualitative, or mixed)
- Written in English
- We excluded:
 - Papers that were not peer-reviewed (e.g., editorials, dissertations)
 - Studies unrelated to healthcare or not clearly focused on MAS
 - Articles that only described theoretical models with no real-world application or evaluation

Search Strategy

A structured search was carried out in four major academic databases:

PubMed, IEEE Xplore, Scopus, and Web of Science.

We used combinations of keywords and Boolean operators, such as:

("multi-agent system" OR "MAS") AND ("healthcare" OR "clinical" OR "medical") AND ("decision support" OR "robotic intervention" OR "critical care")

The search was limited to the last 10 years to ensure relevance to recent advances in AI and MAS technologies deployed in healthcare service delivery. The final search was completed on April 30, 2025.

Study Selection

A total of 150 studies were identified. After removing duplicates (n = 30), the remaining 120 studies were screened by title and abstract. Based on our criteria, 64 full-text articles were reviewed in detail. After applying the inclusion and exclusion filters, 32 studies were selected for the final analysis.

The selection process is summarized in Figure 1 (PRISMA flow diagram) below:




Figure 1: PRISMA flow diagram detailing: records identified (n=150), duplicates removed (n=30), abstracts screened (n=120), abstracts excluded (n=80), full-text articles assessed (n=40), full-text articles excluded (n=8; 5 non-healthcare focus, 3 inadequate reporting), studies included (n=32) [6]

Data Extraction

Two reviewers independently extracted the data using a structured form. Information collected included:

- Year of publication
- Type of healthcare domain (CDSS, robotic system, or critical care)
- MAS design and architecture (e.g., number of agents, type of interaction)
- Evaluation method (e.g., simulations, real-world testing, outcomes measured)
- Performance metrics (accuracy, response time, usability, etc.)
- Ethical or legal considerations mentioned
- Study limitations

Any disagreement between the reviewers was discussed until both reached a consensus.




Figure 2: A visual summary of the key information extracted from each study and the consensus method used to resolve reviewer disagreements.

Quality Assessment

We used the Joanna Briggs Institute (JBI) Critical Appraisal Checklist to assess the quality of each study [7]. This checklist includes questions related to study design, clarity of objectives, data collection, and interpretation of results.

Each study was scored on a scale from 0 to 20, with the following quality levels:

- High quality = 17–20
- Moderate quality = 13–16

- Low quality: < 13

In addition to the JBI criteria, we gave extra attention to:

- How clearly the study explained agent collaboration within the MAS
- Whether patient data privacy and security were addressed
- Whether the MAS was tested in a real-world clinical setting or just in a simulation

Two independent reviewers completed the scoring. If their scores differed by more than two points, the study was reassessed together until they reached agreement.

III. Result

Study Characteristics

Among the 32 studies that met the inclusion criteria:

- 12 focused on Clinical Decision Support Systems (CDSS)
- 11 addressed robotic interventions
- 9 explored MAS applications in critical care

The majority of studies originated from the USA (n = 13), China (n = 8), the UK (n = 4), Germany (n = 5), and India (n = 2). In terms of evaluation methods, 20 were simulation-based, and 12 involved real-world implementation.

MAS in Clinical Decision Support Systems (CDSS)

Multi-agent systems were frequently applied to CDSS to assist with diagnosis, prognosis, and treatment decisions. Architectures commonly used includes Belief-Desire-Intention (BDI), fuzzy logic, and reinforcement learning.

A notable example is PathFinder, a multi-agent system developed for histopathology image analysis. It incorporates triage, navigation, description, and diagnostic agents working collaboratively to classify melanomas. PathFinder outperformed conventional deep learning approaches and even surpassed average pathologist accuracy by 9% [8].

Another implementation involved a modular MAS framework that analyzed patient data from the eICU database using multiple specialized agents (e.g., vital sign analysis, lab interpretation, context-aware inference). This system achieved interpretable and high-performing results in sepsis prediction while maintaining transparency and accountability [9].

MAS in Robotic Interventions

MAS frameworks in robotic healthcare systems supported surgery, rehabilitation, and geriatric care. For instance, MAS-controlled robotic assistants for neurorehabilitation integrated engagement and motor performance feedback to personalize therapy. Yusuf Aydin *et al.*, in 2020 demonstrated how separate agents monitored boredom, focus, and motor response, enhancing outcomes in post-stroke patients [10].

Furthermore, socially assistive robots (SARs) like Paro were managed through multi-agent protocols to adapt interactions based on elderly patients' emotional cues. These implementations showcased improved emotional wellbeing and engagement in dementia care settings [11].

MAS in Critical Care

MAS have also been instrumental in critical care, where patient monitoring, alarm systems, and real-time recommendations are essential. The MATEC (Multi-AI Agent Team Care) system, developed for sepsis detection in rural hospitals, featured coordinated doctor agents, professional-client agents, and a Bayesian-based risk agent. Physicians rated its usefulness highly in early trials [12].

In another example, an ICU monitoring MAS integrated data from physiological sensors and processed it via adaptive agents to generate real-time alerts and ventilator control decisions. This approach provided timely interventions and reduced alarm fatigue [13].

Ethical and Legal Considerations

Only 7 out of 32 studies addressed ethical and legal concerns in depth. Key topics included algorithmic fairness, data privacy, and consent mechanisms. Transparency in agent decision-making was highlighted as essential. Several studies noted regulatory challenges, particularly around compliance with the GDPR, HIPAA, and the proposed EU AI Act of the EU Parliament in 2023 [14]. Liability concerns, especially in autonomous surgical systems, were also discussed.

Comparative Evaluation Matrix

MAS Tool / Domain	Scalability	Clinical Validation	Accuracy	Interpretability	Regulatory Alignment
Oncology CDSS [15-17]	High	Medium	High	Low	Partial
Diabetes CDSS [18,19]	Medium	Low	Medium	Medium	None
Robotic Surgery MAS [20,21]	Medium	High	High	Low	Medium
ICU Monitoring MAS [22,23]	High	High	High	High	High
Parkinson Therapy Bots [24-26]	Low	Medium	Medium	Medium	Partial

Table 1. Comparative Evaluation of MAS Tools across Key Performance Criteria

Descriptive Summary

This evaluation matrix offers a comparative analysis of representative multi-agent systems (MAS) currently implemented or piloted in healthcare. Each system is assessed based on five operational and regulatory criteria:

- **Scalability:** Capacity to extend across facilities or populations while maintaining performance.
- **Clinical Validation:** Extent of real-world testing with peer-reviewed outcomes.
- **Accuracy:** Reliability in delivering correct or expected outputs.
- **Interpretability:** Degree to which system logic is transparent and comprehensible to clinicians.
- **Regulatory Alignment:** Conformity with frameworks such as GDPR, HIPAA, and other medical data standards.

The ICU Monitoring MAS exemplifies a mature system, meeting all five criteria—an outcome likely driven by the high-risk nature of intensive care, where explainability and compliance are mission-critical.

In contrast, the Diabetes CDSS reflects early-stage deployment. Despite moderate accuracy and interpretability, it lacks sufficient validation and regulatory adherence, likely due to limited adoption beyond pilot studies.

The Robotic Surgery MAS demonstrates strong clinical performance but suffers from low interpretability. The complexity of real-time autonomous functions makes transparency a challenge, despite the tool's effectiveness.

Parkinson Therapy Bots represent an emerging niche. These systems monitor engagement, mood, and motor function, offering a personalized approach. Though limited in scale and transparency, early pilots show measurable therapeutic impact, justifying their partial recognition in regulation and interpretability.

Unlike earlier literature, this synthesis reveals a consistent trade-off between interpretability and system complexity; particularly in high-performance MASs that depend on real-time decision-making. This finding suggests a need for design strategies that balance automation with explainability.

Practical Implications

This framework provides actionable insight for multiple stakeholders:

- **Clinicians** can use it to assess readiness and fit for clinical use.
- **Developers** can identify gaps in validation or regulation.
- **Policymakers** may derive benchmarks for certification, especially in areas like transparency and data governance.

IV. Discussion**Interpretation of Results**

This review shows that multi-agent systems (MAS) in healthcare are growing quickly, especially in areas like Clinical Decision Support Systems (CDSS), robotic interventions, and critical care. These systems are showing strong performance in helping healthcare professionals make decisions, improve accuracy, and adapt to different clinical needs. However, their real use in clinical settings is still limited. This is mostly due to a lack of large-scale clinical testing, unclear legal rules, and considerable ethical concerns [27].

Ethical and Legal Implications

Using MAS in healthcare raises important ethical and legal questions:

Human Oversight vs. Machine Autonomy: In places like intensive care units or during robotic surgeries, MAS must support, and not replace clinicians. There must always be a human in the loop to make sure that health service providers stay in control and patient choices are respected [28].

Data Privacy and Security: MAS often work across many systems, using a lot of sensitive data. While laws like HIPAA and GDPR aim to protect patient privacy, many systems still do not use secure methods like encrypted data sharing or federated learning, which would allow learning from data without actually sharing it [29]. Beyond compliance with these regulations, good data governance also means responsible stewardship. This includes clear data management policies, patient consent models adapted for AI, and regular checks to make sure data is used ethically throughout the MAS lifecycle [30].

Transparency and Explainability: Many MAS work like “black boxes” and don’t show how they make decisions. This makes it harder for healthcare service providers and patients to trust them, especially if something goes wrong. Explainable AI (XAI) is important here; because it helps users to understand and trust MAS recommendations [31].

Liability and Accountability: If a MAS makes a mistake that harms a patient, it’s not always clear who is responsible. So many questions come to mind; could it be the developer, the hospital, or the health professional overseeing its use? Legal systems are not yet ready to handle these cases. Regulatory agencies like the FDA, EMA, and new laws like the EU AI Act are starting to develop rules, but more updates are needed [32].

Real-World MAS Deployments

Building on these ethical and legal considerations, it is important to look at real examples of MAS in healthcare:

The Mayo Clinic has developed MAS tools that help monitor patients in the ICU and detect early signs of worsening conditions [33,34].

Stanford Health uses MAS to detect sepsis, combining data from different sources to alert doctors quickly [35,36].

Apollo Hospitals in India are testing MAS for patient rehabilitation, using sensors and robots to help personalize therapy plans [37,38].

Charité Berlin uses MAS in cancer care, helping doctors make more accurate treatment decisions during ongoing trials [39].

These examples show how MAS can help improve care. However, adoption remains limited to a few advanced centres.

Gaps and Limitations

Although MAS show promise, some major gaps remain:

Fewer than 40% of the systems reviewed have been tested in real clinical settings, making it hard to know if they work well outside of labs [40].

Many MAS don’t work well with hospital systems like electronic health records (EHRs), making practical use difficult.

Most MAS are not designed with explainability in mind, which makes it hard for clinicians to trust their results.

Ethical design principles are often missing, and there are no common standards to measure MAS performance and safety.

These problems slow down MAS adoption in hospitals and clinics.

Comparative Framework for MAS Evaluation

This review suggests a simple framework to evaluate MAS based on six key features:

Scalability – Can the system handle large amounts of data and be used in different hospitals?

Speed and Performance – Does it deliver results quickly enough for clinical use?

Computational Efficiency – How well does the system use computing resources (like memory and processors), which affects cost and feasibility?

Clinical Testing – Has it been tested with real patients or in actual clinical environments?

Transparency – Can health professionals and patients understand how it works?

Level of Autonomy – How much does it act on its own, and is human supervision included?

This framework can help healthcare leaders, developers, and policymakers pick MAS that are advanced, safe, fair, and easy to use.

Future Research Directions

To move MAS deployment in health service delivery forward, future research should focus on:

Creating clear ethical and legal guidelines that match how MAS are used in real life [41].

Building MAS that work smoothly with existing clinical software like EHR systems.

Running longer and larger studies in clinical environments to better understand how MAS perform over time.

Including explainable AI tools from the start so healthcare professionals can trust and verify the system's decisions.

Agreeing on common standards to test and compare MAS, so progress can be measured more fairly.

By focusing on these areas, MAS can become more useful, trusted, and ready for wider use in healthcare.

V. Conclusion

Multi-agent AI systems are steadily moving from theory to practice, offering new ways to support clinical decisions, automate care, and respond to patients' changing needs. Their growing use in Clinical Decision Support Systems, robotic interventions, and critical care reflects the broad impact these tools can have across the healthcare sector. Still, real-world success depends on more than technical capability. It requires responsible design, strong clinical validation, safety, and thoughtful alignment with ethical, legal, and regulatory expectations.

Transparent systems that explain their decisions clearly will build trust among healthcare professionals and patients. Ethical frameworks and fair accountability models are especially important in high-stakes areas like surgery and intensive care, where decisions can be life-critical. Similarly, shared evaluation standards will help compare systems, measure their impact in practice, and guide safe, effective adoption.

Looking ahead, meaningful progress will depend on collaboration across disciplines. When researchers, clinicians, developers, and regulators work together, multi-agent AI systems can evolve into reliable and practical tools that truly support better decision-making, safer interventions, and more responsive critical care. Ultimately, these intelligent systems will be key drivers in creating a more personalized, efficient, and patient-centered future for effective healthcare service delivery.

Supplementary Materials:

Supplementary data, including detailed search strategies, the PRISMA flow diagram, and quality assessment checklists, are available at: <https://zenodo.org/records/15116318>

Author Contributions:

Conceptualization, O.P.O. and K.K.; Methodology, O.P.O. and S.A.; Data Curation, O.P.O., N.C.K., A.I., and A.A.; Writing; original draft preparation, O.P.O.; Review and Editing, all authors; Supervision, K.K. All authors have read and agreed to the published version of the manuscript.

Funding:

This research received no external funding.

Institutional Review Board Statement:

Not applicable.

Informed Consent Statement:

Not applicable.

Data Availability Statement:

All data supporting the findings of this review are available within the article and its supplementary materials.

Acknowledgments:

The authors thank their colleagues for their constructive feedback during the preparation of this systematic review.

Conflicts of Interest:

The authors declare no conflict of interest.

References

1. Durfee EH, Rosenschein JS. Distributed problem solving and multi-agent systems: Comparisons and examples. In: Proceedings of the 13th International Distributed Artificial Intelligence Workshop. 1994 Jul;94–104.
2. Kinny D, Georgeff M. Modelling and design of multi-agent systems. In: Müller JP, Wooldridge MJ, Jennings NR, editors. Intelligent Agents III Agent Theories, Architectures, and Languages. ATAL 1996. Lecture Notes in Computer Science, vol 1193. Springer; 1997. p. 56–70. doi:10.1007/BFb0013569.
3. Ribino P, Islam S, Ciampi M, Papastergiou S. Swarm intelligence based multi-agent communication model for securing healthcare ecosystem. In: International Conference on Ubiquitous Computing and Ambient Intelligence. Cham: Springer; 2022 Nov. p. 50–61. doi:10.1007/978-3-031-21333-5_5.

4. Ahmed MI, Spooner B, Isherwood J, Lane M, Orrock E, Dennison A. A systematic review of the barriers to the implementation of artificial intelligence in healthcare. *Cureus*. 2023 Oct 4;15(10):e46454. doi:10.7759/cureus.46454.
5. Li F, Ruijs N, Lu Y. Ethics & AI: A systematic review on ethical concerns and related strategies for designing with AI in healthcare. *AI*. 2022;4(1):28–53.
6. PRISMA. PRISMA 2020 flow diagram [Internet]. Ottawa: PRISMA; 2020 [cited 2025 Mar 31]. Available from: <https://www.prisma-statement.org/prisma-2020-flow-diagram>
7. Santos WMD, Secoli SR, Püschel VADA. The Joanna Briggs Institute approach for systematic reviews. *Rev Lat Am Enfermagem*. 2018;26:e3074. doi:10.1590/1518-8345.2885.3074.
8. Ghezloo F, Seyfioğlu MS, Soraki R, Ikezogwo W, Li B, Vivekanandan T, et al. PathFinder: A Multi-Modal Multi-Agent System for Medical Diagnostic Decision-Making Applied to Histopathology [preprint]. 2025 Feb. Available from: <https://arxiv.org/abs/2502.08916>
9. Chen YJ, Chen CS, Albarqawi A. Reinforcing clinical decision support through multi-agent systems and ethical AI governance [preprint]. 2025 Mar 25. Available from: <https://arxiv.org/abs/2504.03699>
10. Aydin Y, Tokatli O, Patoglu V, Basdogan C. A computational multicriteria optimization approach to controller design for physical human-robot interaction. *IEEE Trans Robot*. 2020 Dec;36(6):1791–804. doi:10.1109/TRO.2020.2998606.
11. Karami V, Yaffe MJ, Gore G, Moon AJ, Rahimi SA. Socially assistive robots for patients with Alzheimer's disease: A scoping review. *Arch Gerontol Geriatr*. 2024;105409. doi:10.1016/j.archger.2024.105409.
12. Cho A, Woo JM, Shi B, Udeshi A, Woo JSH. The application of MATEC (Multi-AI Agent Team Care) framework in sepsis care [preprint]. 2025 Feb 9. Available from: <https://arxiv.org/abs/2503.16433>
13. Li B, Yue L, Nie H, Cao Z, Chai X, Peng B, et al. The effect of intelligent management interventions in intensive care units to reduce false alarms: An integrative review. *Int J Nurs Sci*. 2023 Dec 14;11(1):133–42. doi:10.1016/j.ijnss.2023.12.008.
14. European Parliament. Artificial Intelligence Act: MEPs adopt landmark law [Internet]. 2024 Mar 8 [cited 2025 Jun 10]. Available from: <https://www.europarl.europa.eu/news/en/press-room/20240308IPR19015/artificial-intelligence-act-meps-adopt-landmark-law>
15. Mienye ID, Obaido G, Jere N, Mienye E, Aruleba K, Emmanuel ID, Ogbuokiri B. A survey of explainable artificial intelligence in healthcare: Concepts, applications, and challenges. *Inform Med Unlocked*. 2024;51:101587. doi:10.1016/j.imu.2024.101587.
16. Sharma RK. Enabling scalable and secure healthcare data analytics with cloud-native AI architectures. *Technology (IJRCAIT)*. 2025;8(1).
17. Shapiro MA, Stuhlmiller TJ, Wasserman A, Kramer GA, Federowicz B, Hoos W, et al. AI-augmented clinical decision support in a patient-centric precision oncology registry. *AI Precision Oncol*. 2024;1(1):58–68.
18. Nimri R, Phillip M. Enhancing care in type 1 diabetes with artificial intelligence driven clinical decision support systems. *Horm Res Paediatr*. 2025.
19. Ying Z, Fan Y, Chen C, Liu Y, Tang Q, Chen Z, Chen Y. Real-time AI-assisted insulin titration system for glucose control in patients with type 2 diabetes: A randomized clinical trial. *JAMA Netw Open*. 2025 May;8(5):e258910. doi:10.1001/jamanetworkopen.2025.8910.
20. Jiménez AC, García-Díaz V, Bolaños S. A decentralized framework for multi-agent robotic systems. *Sensors*. 2018 Feb;18(2):417. doi:10.3390/s18020417.
21. Gurusamy KS, Samraj K, Fusai G, Davidson BR. Robot assistant for laparoscopic cholecystectomy. *Cochrane Database Syst Rev*. 2009;(1):CD006578. doi:10.1002/14651858.CD006578.pub2.
22. Rodrigues R, Gonc P, Miranda M, Portela F, Santos M, Neves J, et al. Monitoring intelligent system for the intensive care unit using RFID and multi-agent systems. In: *IEEE International Conference on Industrial Engineering and Engineering Management*; 2012 Dec 10. p. 851–5.
23. Malak JS, Zeraati H, Nayeri FS, Safdari R, Shahraki AD. Neonatal intensive care decision support systems using artificial intelligence techniques: A systematic review. *Artif Intell Rev*. 2019 Dec;52:2685–704.
24. Picelli A, Melotti C, Origano F, Waldner A, Fiaschi A, Santilli V, et al. Robot-assisted gait training in patients with Parkinson disease: A randomized controlled trial. *Neurorehabil Neural Repair*. 2012 May;26(4):353–61. doi:10.1177/1545968311424417.
25. Castelli Gattinara Di Zubiena F, Menna G, Mileti I, Zampogna A, Asci F, Paoloni M, et al. Machine learning and wearable sensors for the early detection of balance disorders in Parkinson's disease. *Sensors*. 2022 Dec 16;22(24):9903.
26. Rodriguez F, Krauss P, Kluckert J, Ryser F, Stieglitz L, Baumann C, et al. Continuous and unconstrained tremor monitoring in Parkinson's disease using supervised machine learning and wearable sensors. *Parkinsons Dis*. 2024;2024:5787563.
27. World Health Organization. Ethics and governance of artificial intelligence for health: WHO guidance. Geneva: WHO; 2021. ISBN 978-92-4-002920-0 (electronic), 978-92-4-002921-7 (print). Available from: <https://creativecommons.org/licenses/by-nc-sa/3.0/igo>
28. Fosch Villaronga E, Khanna P, Drukarch HG, Custers B. A human in the loop in surgery automation. *Nat Mach Intell*. 2021 May;3(5):1–2. doi:10.1038/s42256-021-00349-4.

29. Pati S, Kumar S, Varma A, Edwards B, Lu C, Qu L, et al. Privacy preservation for federated learning in health care. *Patterns* (N Y). 2024;5(3):100974. doi:10.1016/j.patter.2024.100974.
30. Solanki P, Grundy J, Hussain W. Operationalising ethics in artificial intelligence for healthcare: A framework for AI developers. *AI Ethics*. 2023;3:223–40. doi:10.1007/s43681-022-00195-z.
31. Saarela M, Podgorelec V. Recent applications of explainable AI (XAI): A systematic literature review. *Appl Sci*. 2024;14(19):8884. doi:10.3390/app14198884.
32. Fraser AG, Biasin E, Bijmens B, Bruining N, Caiani EG, Cobbaert K, et al. Artificial intelligence in medical device software and high-risk medical devices – a review of definitions, expert recommendations and regulatory initiatives. *Expert Rev Med Devices*. 2023;20(6):467–91. doi:10.1080/17434440.2023.2184685.
33. Romero-Brufau S, Whitford D, Johnson MG, Hickman J, Morlan BW, Therneau T, et al. Using machine learning to improve the accuracy of patient deterioration predictions: Mayo Clinic Early Warning Score (MC-EWS). *J Am Med Inform Assoc*. 2021 Jun 1;28(6):1207–15.
34. Romero-Brufau S, Whitford D, Johnson MG, Hickman J, Morlan BW, Therneau T, et al. Using machine learning to improve the accuracy of patient deterioration predictions: Mayo Clinic Early Warning Score (MC-EWS). *J Am Med Inform Assoc*. 2021 Jun 1;28(6):1207–15. doi:10.1093/jamia/ocab012.
35. Wong A, Cao J, Lyons PG, Dutta S, Major VJ, Otles E, Singh K. Quantification of sepsis model alerts in 24 US hospitals before and during the COVID-19 pandemic. *JAMA Netw Open*. 2021 Nov 1;4(11):e2135286.
36. Slack G. A better, faster sepsis test that can save lives. *Innovation & Technology*. 2025 Feb 5. Available from: <https://med.stanford.edu/news/insights/2025/02/sepsis-infection-new-test-genetic.html>
37. Gobinath A, Rajeswari P, Suresh KN, Anandan M. Robotics in real-time applications in healthcare systems. In: *AI and IoT Technology and Applications for Smart Healthcare Systems*. Auerbach Publications; 2024 May 15. p. 262–74.
38. Hospital Management Asia. AI-powered patient monitoring system at Apollo Hospitals cut Code Blues by 80%. 2025 Feb 6. Available from: <https://www.hospitalmanagementasia.com/patient-safety/ai-powered-patient-monitoring-system-at-apollo-hospitals-cut-code-blues-by-80/>
39. Verlingue L, Boyer C, Olgiaiti L, Brutti M, Morel D, Blay JY. Artificial intelligence in oncology: ensuring safe and effective integration of language models in clinical practice. [Internet]. 2024 Nov [cited 2025 Jun 10]. Available from: [Publisher link unavailable].
40. He J, Baxter SL, Xu J, Zhou X, Zhang K. The practical implementation of artificial intelligence technologies in medicine. *Nat Med*. 2019 Jan;25(1):30–6.
41. Morley J, Machado CC, Burr C, Cows J, Joshi I, Taddeo M, Floridi L. The ethics of AI in health care: A mapping review. *Soc Sci Med*. 2020;260:113172.