

Multiagent AI Systems in Health Care: Envisioning Next-Generation Intelligence

Andrew A. Borkowski, MD^a; Alon Ben-Ari, MD, MSc^b

Background: Limited staff, rising costs, and regulatory oversight, coupled with the need to achieve clinical endpoints and improve access to care, has made scaling health care operations challenging. This article explores the emerging paradigm of multiagent artificial intelligence (AI) systems in health care, which represent a significant leap beyond traditional large language models.

Observations: This analysis reviews the potential of multiagent AI systems to revolutionize patient care, streamline administrative processes, and support complex clinical decision-making. It describes a hypothetical sepsis management system comprising 7 specialized AI agents, with each agent handling specific aspects of patient care from data collection and diagnosis to treatment recommendations and resource management. Additional applications in chronic disease management and hospital patient flow optimization are also examined. The technical

implementation of these systems is discussed, including the use of advanced large language models, interagent quality control measures, guardrail implementation, self-reflection mechanisms, integration with electronic health records, and the importance of explainable AI in ensuring decision transparency. Potential benefits include enhanced diagnostic accuracy and personalized treatment plans. Challenges remain related to data quality assurance, workflow integration, and ethical considerations. Future directions for AI include the integration of internet-enabled devices and the development of more sophisticated natural language interfaces.

Conclusions: This article underscores the transformative potential of multiagent AI systems in health care while emphasizing the importance of rigorous validation, ethical oversight, and a patient-centered approach in their development and implementation.

Author affiliations can be found at the end of this article.

Correspondence:

Andrew Borkowski
(andrew.borkowski@va.gov)

Fed Pract. 2025;42(5).
Published online May 14.
doi:10.12788/fp.0589

Artificial intelligence (AI) is rapidly evolving, with large language models (LLMs) marking a significant milestone in processing and generating human-like responses to natural language prompts. However, this advancement only signals the beginning of a more profound transformation in AI capabilities. The development of AI agents represents a new paradigm at the forefront of this evolution.

BACKGROUND

AI agents represent a leap forward from traditional LLM applications. While definitions may vary slightly among technology developers, the core concept remains: these agents are autonomous software entities designed to interact with their environment, make independent decisions, and execute tasks based on predefined goals.¹⁻³ What sets AI agents apart is their combination of sophisticated components within structured architectures. At their core, AI agents incorporate an LLM for response generation, which is augmented by a suite of tools to optimize workflow and complete tasks, memory capabilities for personalized interactions, and autonomous reasoning. This combination allows AI agents to plan, create subtasks, gather information, and learn iteratively from

their own experiences or other AI agents.

The true potential of this technology becomes apparent when multiple AI agents collaborate within multiagent AI systems. This concept introduces a new level of flexibility and capability in tackling complex tasks. Autogen, CrewAI, and LangChain offer various agent network configurations, including hierarchical, sequential, conditional, or even parallel task execution.⁴⁻⁶ This adaptability opens up a world of possibilities across various industries, but perhaps nowhere is the potential impact more exciting and profound than in health care.

AI agents in health care present an opportunity to revolutionize patient care, streamline administrative processes, and support complex clinical decision-making. This review examines 3 scenarios that illustrate the impact of AI agents in health care: a hypothetical sepsis management system, chronic disease management, and hospital patient flow optimization. This article will provide a detailed look at the technical implementation challenges, including the integration with existing health care IT systems, data privacy considerations, and the crucial role of explainable AI in maintaining trust and transparency.

It is challenging to implement AI agents in

health care. Concerns include ensuring data quality and mitigating bias, seamlessly integrating these systems into existing clinical workflows, and navigating the complex ethical considerations that arise when deploying autonomous systems in health care. The integration with Internet of Things (IoT) devices for real-time patient data monitoring and the development of more sophisticated natural language interfaces to enhance future human-AI collaboration.

The adoption of AI agents in health care is only beginning, and it promises to be transformative. As AI continues to evolve, a comprehensive understanding of its applications, limitations, and ethical considerations is essential. This report provides a comprehensive overview of the current state, potential applications, and future directions of AI agents in health care, offering insights valuable to researchers, clinicians, and policymakers.

MULTIAGENT AI ARCHITECTURE

Sepsis Management

Despite advancements in broad-spectrum antibiotics, imaging, and life support systems, mortality rates associated with sepsis remain high. The complexity of optimizing care in clinical settings has hindered progress in managing sepsis. Previous attempts to develop predictive sepsis models have proven challenging.⁷ This report proposes a multiagent AI system designed to enhance comprehensive patient monitoring and care through coordinated AI-driven interventions.

Data Collection and Integration Agent. Powered by a controlled vocabulary to specify all data, the primary function for the data collection and integration agent is to clean, transform, and organize patient data from structured and unstructured sources. This agent prepares succinct summaries of consultant notes and formats data for human and machine consumption. All numerical data are presented graphically, including relevant historical data trends. The agent also digitally captures all orders in a structured format using a specified controlled vocabulary. This structured data feed supports the output of other agents, including documentation, treatment planning, and risk stratification, while also supplying the data structures for future training.

Diagnostic Agent. Critical illness is characterized by multiple abnormalities across a wide array of tests, ranging from plain chest X-ray, computed tomography (CT), blood cell composition, plasma chemistry, and microscopic evaluation of specimens. Additionally, life support parameters provide insights into disease severity and can inform management recommendations. These data offer a wide array of visual and numerical data to be used as input for computation, recommendation, and further training. For example, to evaluate fluid overload on chest X-rays or tissue histopathology slides, an AI agent can leverage deep learning models such as convolutional neural networks and vision transformers to analyze images like radiographs and histopathology slides.^{8,9} Recurrent neural networks or transformer models process sequential data like time-series vital signs. The agent also implements ensemble methods that combine multiple machine learning algorithms to enhance diagnostic accuracy.

Risk Stratification Agent. This assesses severity and predicts potential outcomes. Morbidity and mortality risks are calculated using an established scoring system and individualized based on the history of other agents' conditional patients. These are presented graphically, with major risk factors highlighted for explainability.

Treatment Recommendation Agent. Using a reinforcement learning framework supplemented by up-to-date clinical guidelines, this system leverages historical data structured with standardized vocabulary to analyze patients with similar clinical features. Training is also conducted on the patient's physiological data. All recommendations are presented via a dedicated user interface in a readable format, along with recommendations for editable, orderable items, references, and full-text snippets from previous research. Stop rules end computing if confidence in recommendations is too broad or no clear pathway can be computed with certainty, prompting human mitigation.

Resource Management Agent. This agent coordinates hospital resources using constraint programming techniques for optimal resource allocation, uses queueing theory models to predict and manage patient flow, and implements genetic algorithms for com-

plex scheduling problems.^{10,11}

Monitoring and Alert Agent. By tracking patients' progress and alerting staff to changes, this agent uses anomaly detection algorithms to identify unusual patterns in patient data and implement time-series forecasting models, such as autoregressive integrated moving average and prophet, to predict future patient states. The agent also uses stream-processing techniques for real-time data analysis.^{12,13}

Documentation and Reporting Agent. This agent maintains comprehensive medical records and generates reports. It employs advanced natural language processing techniques for automated report generation, uses advanced LLMs fine-tuned on medical corpora for narrative creation, and implements information-retrieval techniques to efficiently query patient records.

CLINICAL CASE STUDIES

To illustrate the functionality of a multi-agent system, this report examines its application for managing sepsis. The data collection and integration agent continuously aggregates patient data from various sources, normalizing and timestamping it for consistent processing. The diagnostic agent analyzes this integrated data in real time, applying sepsis criteria and utilizing a deep learning model trained on a large sepsis dataset to detect subtle patterns.

The risk stratification agent calculates severity scores, such as the Sepsis-related Organ Failure Assessment (SOFA), quick SOFA (qSOFA), and Acute Physiology and Chronic Health Evaluation II, upon detecting a possible sepsis case.¹⁴ It predicts the likelihood of specific outcomes and estimates the potential trajectory of the patient's condition for the next 24 to 48 hours. Based on this assessment, the treatment recommendation agent suggests an initial treatment plan, including appropriate antibiotics, fluid resuscitation protocols, and vasopressor recommendations, recommendations when indicated.

Concurrently, the resource management agent checks the availability of necessary resources and prioritizes allocation based on the severity. The monitoring agent tracks the patient's response to interventions in real time, alerting the care team to any con-

cerning changes or lack of expected improvement. Throughout this process, the documentation agent ensures that all actions, responses, and outcomes are meticulously recorded in a structured format and generates real-time updates for the patient's electronic health record (EHR) and preparing summary reports for handoffs between care teams.

Administrative Workflow Support

Modern health care operations are resource-intensive, requiring coordination of advanced imaging, procedures, laboratory testing, and professional consultations.¹⁵ AI-powered health care administrative workflow systems are revolutionizing how medical facilities coordinate patient care. For patients with chronic cough, these systems seamlessly integrate scheduling, imaging, diagnostics, and follow-up care into a cohesive process that reduces administrative burden while improving patient outcomes. Through an intuitive interface and automated assistance, health care practitioners (HCPs) can track patient progress from initial consultation through diagnosis and treatment.

The process begins when an HCP enters a patient into the system, which triggers an automated CT scan scheduling system. The system considers factors like urgency, facility availability, and patient preferences to suggest optimal appointment times. Once imaging is complete, AI agents analyze the radiology reports, extract key findings, and generate structured summaries that highlight critical information such as "mild bronchial wall thickening with patchy ground-glass opacities" or "findings consistent with chronic bronchitis."

Based on these findings, the system automatically generates evidence-based recommendations for follow-up care, such as pulmonology consultations or follow-up imaging in 3 months. These recommendations are presented to the ordering clinician, along with suggested appointment slots for specialist consultations. The system then manages the coordination of multiple appointments, ensuring each step in the patient's care plan is properly sequenced and scheduled.

The entire process is monitored through a comprehensive dashboard that provides real-

time updates on patient status, appointment schedules, and clinical recommendations. HCPs can track which patients require immediate attention, view upcoming appointments, and monitor the progress of ongoing care plans.

Multiagent AI Operation Optimization

Hospitals are complex entities that must function at different scales and respond in an agile, timely manner at all hours, deploying staff at various positions.¹⁶ A system of AI agents can receive signals from sensors monitoring foot traffic in the emergency department and trauma unit, as well as the availability of operating room staff, equipment, and intensive care unit beds. Smart sensors enable this monitoring through IoT networks. These networks benefit from advances in adaptive and consensus networking algorithms, along with recent advances in bioengineering and biocomputing.¹⁷

For example, in the case of imaging for suspected abdominal obstruction, an AI agent tasked with scheduling CTs could time the patient's arrival based on acuity. Another AI agent could alert staff transporting the patient to the CT appointment, with the next location contingent on a clinical decision to proceed to the operating room. Yet another AI agent could summarize radiology interpretations and alert the surgery and anesthesia teams to a potential case, while others could notify operating room staff of equipment needs or reserve a bed. In this paradigm, AI agents facilitate more precise and timely communication between multiple staff members.

TECHNICAL IMPLEMENTATION

Large Language Models

Each agent uses a different LLM optimized for its specific task. For example, the diagnostic agent uses an LLM pre-trained on a large corpus of biomedical literature and fine-tuned on a dataset of confirmed sepsis cases and their presentations.¹⁸ It implements few-shot learning techniques to adapt to rare or atypical presentations. The treatment recommendation agent also uses an LLM, employing a retrieval-augmented generation approach to access the latest clinical guidelines during inference. The docu-

mentation agent uses another advanced language model, fine-tuned on a large corpus of high-quality medical documentation, implementing controlled text generation techniques and utilizing a separate smaller model for real-time error checking and correction.

Interagent Quality Control

Agents learn from their own experience and the experience of other agents. They are equipped with user-defined rule-based and model-based systems for quality assurance, with clear stopping rules for human involvement and mitigation.

Sophisticated quality control measures bolster the system's reliability, including ensemble techniques for result comparison, redundancy for critical tasks, and automatic human review for disagreements above a certain threshold. Each agent provides a calibrated confidence score with its output, used to weigh inputs in downstream tasks and trigger additional checks for low-confidence outputs.

A dedicated quality control agent monitors output from all agents, employing both supervised and unsupervised anomaly detection techniques. Feedback loops allow agents to evaluate the quality and utility of information received from other agents. The system implements a multiarmed bandit approach to dynamically adjust the influence of different agents based on their performance and periodically retrains agent models using federated learning techniques.¹⁹

Electronic Health Record Integration

Seamless EHR integration is crucial for practical implementation. The system has secure application programming interface access to various EHR platforms, implements OAuth 2.0 for authentication, and use HTTPS with perfect forward secrecy for all communications.²⁰ It works with HL7 FHIR to ensure interoperability and uses SNOMED CT for clinical terminology to ensure semantic interoperability across different EHRs.^{21,22}

The system implements a multilevel approval system for write-backs to EHRs, with different thresholds based on the information's criticality. It uses digital signatures to ensure the integrity and nonrepudiation of

AI-generated entries and implements blockchain technology to create an immutable and distributed ledger of all AI system actions.²³

Decision Transparency

To ensure transparency in decision-making processes, the system applies techniques (eg, local interpretable model-agnostic explanations and Shapley additive explanations) to provide insights into agent decision-making processes.²⁴⁻²⁶ It provides customized visualizations for different stakeholders and allows users to explore alternative decision paths through what-if scenario modeling.²⁷

The system provides calibrated confidence indicators for each recommendation or decision, implementing a novel confidence calibration agent that continuously monitors and adjusts confidence scores based on observed outcomes.

Continuous Learning and Adaptation

The system employs several techniques to remain current with evolving medical knowledge. Federated learning includes information from diverse datasets across multiple institutions without compromising patient privacy.²⁸ A/B testing is used to safely deploy and compare new agent versions in controlled settings, implementing multiarmed bandit algorithms to efficiently explore new models while minimizing potential negative impacts. Human-in-the-loop learning and active learning techniques are used to incorporate feedback from HCPs and efficiently solicit expert input on the most informative data.²⁹

CLINICAL IMPLICATIONS

The implementation of multiagent AI systems in health care has several potential benefits: enhanced diagnostic accuracy, personalized treatment, improved efficiency, continuous monitoring, and resource optimization. A recent review of AI sepsis predictive models exhibited superior results to standard clinical scoring methods like qSOFA.³⁰ In oncology, such systems can result in more tailored treatments, enhancing outcomes.³¹ The implementation of an ambient dictation system can improve workflow and prevent HCP burnout.³²

ETHICAL CONSIDERATIONS AND AI OVERSIGHT

Integrating AI agents into health care raises significant ethical considerations that must be carefully addressed to ensure equitable and effective care delivery. One primary concern involves cultural and linguistic competency, as AI systems may struggle with cultural nuances, idioms, and context-specific communication patterns. This becomes particularly challenging in regions with diverse ethnic populations or immigrant communities, where medical terminology may not have direct translations and cultural beliefs significantly influence health care decisions. AI systems also may inherit and amplify existing biases in health care delivery, whether through HCP bias reflected in training data, patient bias affecting acceptance of AI-assisted care, or demographic underrepresentation during system development.

AI agents present unique opportunities for improving health care access and outcomes through community engagement, though such initiatives require thoughtful implementation. Predictive analytics can identify high-risk individuals within communities who may benefit from preventive care, while analysis of social determinants of health can enable more targeted interventions. However, these capabilities must be balanced with privacy concerns and the risk of surveillance, particularly in communities that distrust health care institutions. The potential for AI to bridge health care gaps must be weighed against the need to maintain cultural sensitivity and community trust.

The governance and oversight of health care AI systems requires a multistakeholder approach with clear lines of responsibility and accountability. This includes involvement from government health care agencies, professional medical associations, ethics boards, and independent auditors, all working together to establish and enforce standards while monitoring system performance and addressing potential biases. Health care organizations must maintain transparent policies about AI use, implement regular monitoring and evaluation protocols, and establish precise mechanisms for patient feedback and grievance resolution. Ongoing assessment and adjustment of these systems, informed

by community feedback and outcomes data, will be crucial for their ethical implementation, ensuring that AI agents complement, rather than replace, human judgment and cultural sensitivity.

FUTURE DIRECTIONS

Despite the potential benefits, implementing multiagent AI systems in health care faces significant challenges that require careful consideration. Beyond the fundamental issues such as data quality and bias mitigation, health care organizations struggle with fragmented systems, inconsistent data formats, and varying quality. Technical infrastructure requirements are substantial, particularly in rural or underserved areas that lack robust networks and cybersecurity. HCPs already face significant cognitive load and time pressures, making integrating AI agents into existing workflows particularly challenging. There is also the critical issue of transparency and interpretability, as health care decisions require clear reasoning and accountability that many black-box AI systems struggle to provide.

The legal landscape introduces another layer of complexity, particularly regarding liability, consent, and privacy questions. When AI agents contribute to medical decisions, establishing clear lines of responsibility becomes crucial. There are also serious concerns about algorithmic fairness and the potential for AI systems to perpetuate or amplify existing inequities. The cost of implementation remains a significant barrier, requiring substantial investment in technology, training, and ongoing maintenance while ensuring resources are not diverted from direct patient care. Moreover, HCPs may resist adoption due to concerns about job security, loss of autonomy, or skepticism about AI capabilities while paradoxically facing risks of overreliance on AI systems that could lead to the degradation of human clinical skills.

Addressing these challenges requires a multifaceted approach that combines technical solutions with organizational and policy changes. Health care organizations must implement rigorous data validation processes and interoperability standards while developing hybrid models that balance sophisticated AI capabilities with interpretable

techniques. Extensive research and iterative design processes, with direct input from HCPs, are essential for successful integration. Establishing independent ethics boards to oversee system development and deployment, conducting multicenter randomized controlled trials, and creating clear regulatory frameworks will ensure safe and effective implementation. Success will ultimately depend on ongoing collaboration between technology developers, HCPs, policymakers, and patients, maintaining a steady focus on improving patient care and outcomes while carefully navigating the complex challenges of AI integration in health care.³³⁻³⁵

As multiagent AI systems in health care evolve, several exciting directions emerge. These include the integration of IoT and wearable devices, the development of more sophisticated natural language interfaces, and applying these systems to predictive maintenance of medical equipment.

CONCLUSIONS

The advent of multiagent AI systems in health care represents a paradigm shift in the approach to patient care, clinical decision making, and health care management. While these systems offer immense potential to transform health care delivery, their development and implementation must be guided by rigorous scientific validation, ethical considerations, and a patient-centered approach. The ultimate goal remains clear: harnessing the power of AI to improve patient outcomes, enhance the efficiency of health care delivery, and ultimately advance the health and well-being of patients.

Author affiliations

^aVeterans Affairs Sunshine Healthcare Network, Tampa, Florida
^bVeterans Affairs Northern California Health Care System, Sacramento

Author disclosures

The authors report no actual or potential conflicts of interest with regard to this article.

Disclaimer

The opinions expressed herein are those of the authors and do not necessarily reflect those of *Federal Practitioner*, Frontline Medical Communications Inc., the US Government, or any of its agencies.

References

1. Amazon Web Services, Inc. What are AI agents? Agents in artificial intelligence explained. Accessed April 7, 2025. <https://aws.amazon.com/what-is/ai-agents/>
2. Gutowska A. What are AI agents? IBM. Accessed April 7, 2025. <https://www.ibm.com/think/topics/ai-agents>

3. Agent AI. Microsoft Research. Accessed April 7, 2025. <https://www.microsoft.com/en-us/research/project/agent-ai>
4. Microsoft. AutoGen. Accessed April 7, 2025. <https://microsoft.github.io/autogen/>
5. Crew AI. The Leading Multi-Agent Platform. CrewAI. Accessed April 7, 2025. <https://www.crewai.com/>
6. LangChain. Accessed April 7, 2025. <https://www.langchain.com/>
7. Wong A, Otles E, Donnelly JP, et al. External validation of a widely implemented proprietary sepsis prediction model in hospitalized patients. *JAMA Intern Med.* 2021;181(8):1065-1070. doi:10.1001/jamainternmed.2021.2626
8. Willemink MJ, Roth HR, Sandfort V. Toward foundational deep learning models for medical imaging in the new era of transformer networks. *Radiol Artif Intell.* 2022;4(6):e210284. doi:10.1148/ryai.210284
9. Waqas A, Bui MM, Glassy EF, et al. Revolutionizing digital pathology with the power of generative artificial intelligence and foundation models. *Lab Invest.* 2023;103(11):100255. doi:10.1016/j.labinv.2023.100255
10. Moreno-Carrillo A, Arenas LMÁ, Fonseca JA, Caicedo CA, Tovar SV, Muñoz-Velandia OM. Application of queuing theory to optimize the triage process in a tertiary emergency care ("ER") department. *J Emerg Trauma Shock.* 2019;12(4):268-273. doi:10.4103/JETS.JETS_42_19
11. Pongcharoen P, Hicks C, Braiden PM, Stewardson DJ. Determining optimum genetic algorithm parameters for scheduling the manufacturing and assembly of complex products. *Int J Prod Econ.* 2002;78(3):311-322. doi:10.1016/S0925-5273(02)00104-4
12. Sardar I, Akbar MA, Leiva V, Alsanad A, Mishra P. Machine learning and automatic ARIMA/Prophet models-based forecasting of COVID-19: methodology, evaluation, and case study in SAARC countries. *Stoch Environ Res Risk Assess.* 2023;37(1):345-359. doi:10.1007/s00477-022-02307-x
13. Samosir J, Indrawan-Santiago M, Haghighi PD. An evaluation of data stream processing systems for data driven applications. *Procedia Comput Sci.* 2016;80:439-449. doi:10.1016/j.procs.2016.05.322
14. Asmarawati TP, Suryantoro SD, Rosyid AN, et al. Predictive value of sequential organ failure assessment, quick sequential organ failure assessment, acute physiology and chronic health evaluation II, and new early warning signs scores estimate mortality of COVID-19 patients requiring intensive care unit. *Indian J Crit Care Med.* 2022;26(4):466-473. doi:10.5005/jp-journals-10071-24170
15. Khan S, Vandermorris A, Shepherd J, et al. Embracing uncertainty, managing complexity: applying complexity thinking principles to transformation efforts in healthcare systems. *BMC Health Serv Res.* 2018;18(1):192. doi:10.1186/s12913-018-2994-0
16. Plsek PE, Greenhalgh T. The challenge of complexity in health care. *BMJ.* 2001;323(7313):625-628. doi:10.1136/bmj.323.7313.625
17. Kouchaki S, Ding X, Sanei S. AI- and IoT-enabled solutions for healthcare. *Sensors.* 2024;24(8):2607. doi:10.3390/s24082607
18. Saab K, Tu T, Weng WH, et al. Capabilities of Gemini Models in Medicine. *arXiv.* doi:10.48550/arXiv.2404.18416
19. Villar SS, Bowden J, Wason J. Multi-armed bandit models for the optimal design of clinical trials: benefits and challenges. *Stat Sci.* 2015;30(2):199-215. doi:10.1214/14-ST504
20. Auth0. What is OAuth 2.0. Accessed April 7, 2025. <https://auth0.com/intro-to-iam/what-is-oauth-2>
21. HL7. Welcome to FHIR. Updated March 26, 2025. Accessed April 7, 2025. <https://www.hl7.org/fhir/>
22. SNOMED International. Accessed April 7, 2025. <https://www.snomed.org>
23. Hasselgren A, Kravetska K, Gligoroski D, Pedersen SA, Faxvaag A. Blockchain in healthcare and health sciences—a scoping review. *Int J Med Inf.* 2020;134:104040. doi:10.1016/j.ijmedinf.2019.104040
24. Ribeiro MT, Singh S, Guestrin C. "Why Should I Trust You?": Explaining the predictions of any classifier. In: *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining.* 2016:1135-1144. doi:10.1145/2939672.2939778
25. Ekanayake IU, Meddage DPP, Rathnayake U. A novel approach to explain the black-box nature of machine learning in compressive strength predictions of concrete using Shapley additive explanations (SHAP). *Case Stud Constr Mater.* 2022;16:e01059. doi:10.1016/j.cscm.2022.e01059
26. Alabi RO, Elmusrati M, Leivo I, Almangush A, Mäkitie AA. Machine learning explainability in nasopharyngeal cancer survival using LIME and SHAP. *Sci Rep.* 2023;13(1):8984. doi:10.1038/s41598-023-35795-0
27. Otto E, Culakova E, Meng S, et al. Overview of sankey flow diagrams: focusing on symptom trajectories in older adults with advanced cancer. *J Geriatr Oncol.* 2022;13(5):742-746. doi:10.1016/j.jgo.2021.12.017
28. Fereidooni H, Marchal S, Miettinen M, et al. SAFElearn: secure aggregation for private federated learning. In: *2021 IEEE security and privacy workshops (SPW).* 2021:56-62. doi:10.1109/SPW53761.2021.00017
29. Linton DL, Pangle WM, Wyatt KH, Powell KN, Sherwood RE. Identifying key features of effective active learning: the effects of writing and peer discussion. *Life Sci Educ.* 2014;13(3):469-477. doi:10.1187/cbe.13-12-0242
30. Yang HS. Machine learning for sepsis prediction: prospects and challenges. *Clin Chem.* 2024;70(3):465-467. doi:10.1093/clinchem/hvae006
31. Liao J, Li X, Gan Y, et al. Artificial intelligence assists precision medicine in cancer treatment. *Front Oncol.* 2023;12. doi:10.3389/fonc.2022.998222
32. Tierney AA, Gayre G, Hoberman B, et al. Ambient artificial intelligence scribes to alleviate the burden of clinical documentation. *NEJM Catal.* 2024;5(3):CAT.23.0404. doi:10.1056/CAT.23.0404
33. Borkowski AA, Jakey CE, Thomas LB, Viswanadhan N, Mastorides SM. Establishing a hospital artificial intelligence committee to improve patient care. *Fed Pract.* 2022;39(8):334-336. doi:10.12788/fp.0299
34. Isaacs DB, Borkowski AA. Implementing trustworthy AI in VA high reliability health care organizations. *Fed Pract.* 2024;41(2):40-43. doi:10.12788/fp.0454
35. Han R, Acosta JN, Shakeri Z, Ioannidis JPA, Topol EJ, Rajpurkar P. Randomized controlled trials evaluating artificial intelligence in clinical practice: a scoping review. *Lancet Digit Health.* 2024;6(5):e367-e373. doi:10.1016/S2589-7500(24)00047-5

