By Michael E. Matheny, Jennifer C. Goldsack, Suchi Saria, Nigam H. Shah, Jacqueline Gerhart, I. Glenn Cohen, W. Nicholson Price II, Bakul Patel, Philip R. O. Payne, Peter J. Embí, Brian Anderson, and Eric Horvitz

DOI: 10.1377/hlthaff.2024.01003
HEALTH AFFAIRS 44,
NO. 2 (2025): 163–170


COMMENTARY

# Artificial Intelligence In Health And Health Care: Priorities For Action

**ABSTRACT The field of artificial intelligence (AI) has entered a new cycle of intense opportunity, fueled by advances in deep learning, including generative AI. Applications of recent advances affect many aspects of everyday life, yet nowhere is it more important to use this technology safely, effectively, and equitably than in health and health care. Here, as part of the National Academy of Medicine's Vital Directions for Health and Health Care: Priorities for 2025 initiative, which is designed to provide guidance on pressing health care issues for the incoming presidential administration, we describe the steps needed to achieve these goals. We focus on four strategic areas: ensuring safe, effective, and trustworthy use of AI; promotion and development of an AI-competent health care workforce; investing in AI research to support the science, practice, and delivery of health and health care; and promotion of policies and procedures to clarify AI liability and responsibilities.**

**Michael E. Matheny** (michael.matheny@vumc.org), Vanderbilt University and Veterans Affairs Tennessee Valley Healthcare System, Nashville, Tennessee.

**Jennifer C. Goldsack**, Digital Medicine Society, Boston, Massachusetts.

**Suchi Saria**, Johns Hopkins University, Baltimore, Maryland.

**Nigam H. Shah**, Stanford University, Palo Alto, California.

**Jacqueline Gerhart**, Epic, Verona, Wisconsin.

**I. Glenn Cohen**, Harvard University, Boston, Massachusetts.

**W. Nicholson Price II**, University of Michigan, Ann Arbor, Michigan.

**Bakul Patel**, Google, Washington, D.C.

**Philip R. O. Payne**, Washington University in St. Louis, St. Louis, Missouri.

**Peter J. Embí**, Vanderbilt University.

**Brian Anderson**, Coalition for Health AI, Boston, Massachusetts.

**Eric Horvitz**, Microsoft, Redmond, Washington.

Efforts to incorporate artificial intelligence (AI) technologies into health care have been under way for decades. Recently, rapid evolution in the development of AI technologies has provided new possibilities for health care while also challenging health care organizations, legislators, policy makers, and regulators to provide guidance to support reliable and safe, yet innovative, systems. US policy makers have moved quickly to issue relevant laws, policies, and regulatory guidelines. Consensus organizations have also sought to interpret and identify gaps and expand recommendations for AI best practices (see online appendix exhibit 1).[1] These reports and guidelines collectively emphasize the importance of transparency, trustworthiness, fairness and equity, data interoperability and accessibility, and patient safety. They highlight the need for inclusive collaboration, ongoing safety assessment, and governance processes and infrastructure to manage risks associated with AI implementation. These reports stress the importance of addressing disparities and inequities in health care, cautioning against the potential for unintended consequences and exacerbation of disparities as AI systems are implemented.

This article is part of the National Academy of Medicine's initiative, Vital Directions for Health and Health Care: Priorities for 2025, which critically assesses public health issues to provide guidance for the incoming US presidential administration. We first describe historical context regarding the advancements and breakthroughs in AI that have led to this critical point. Next, we identify four key policy-related domains that are critical to enabling the use of AI to promote the health of all Americans.

## Historical Context

Foundational research in health care AI began in the late 1950s in two areas: symbolic representa-

tions using logic; and statistical methods, including inference with probabilities and with a predecessor to today's neural network models, called perceptrons.[2] These approaches evolved over decades as experts explored multiple paradigms of reasoning and representation, including methods that generate advice via chaining of rules acquired from experts and probabilistic approaches, such as Bayesian networks.

More recent machine learning methods construct inference systems entirely from health care data sets to enable the development of diagnostic, predictive, and generative systems. This capacity has been enabled by sweeping expansions in electronic data acquisition and computational capacities.[2] However, despite excitement and opportunity, AI technologies have achieved only limited diffusion and impact in health care practice.

Historically, the most clinically impactful AI technologies have not been commonly recognized as AI by the health care community. Widely used scoring algorithms such as the ten-year cardiovascular risk calculator; the Length of Stay, Acuity, Comorbidities, and Recent Emergency Department Use (LACE) index for readmission prediction;[3] the Ottawa Ankle Rules for fracture assessment;[4] the Acute Physiology and Chronic Health Evaluation (APACHE) score for intensive care unit mortality;[5] and the $CHA_2DS_2$-VASc score[6] for stroke risk in atrial fibrillation are examples of tools developed from early discriminative AI models and adapted so that they could be calculated manually in practice. The clinical efficacy of several of these scores has been compared with power delivered by machine learning models.[7]

A new wave of promising tools is coming to health care with advances in deep neural networks. Deep neural networks provide different types of inferential tools with new capacity for classification, diagnosis, prediction, and generative models. They can produce rich simulations; engage in consultative dialogues; and perform language-centric tasks that are being incorporated into administrative and clinical workflows, including summarization, report writing, and accurate transcription from voice input. In clinical workflows, deep neural network–enabled tools and larger systems harnessing them as components show great promise for assisting with digital imaging processing, diagnostic reasoning, and population health management.[8] These tools can boost the timeliness and accuracy of diagnosis and therapy and minimize errors and delays linked to adverse outcomes.

## Four Key Domains For AI In Health And Health Care

As noted, the rapid evolution of AI technologies and their potential application in health and health care has raised important opportunities and challenges for legislators, policy makers, and regulators. Here, we identify and describe four key policy-related domains, each with multiple elements, that are critical to enabling the use of AI to promote the health of all Americans. The domains are ensuring the safe, effective, and trustworthy use of AI; promoting the development of an AI-competent workforce; promoting research on AI in health and health care; and clarifying responsibility and liability in the use of AI. A summary of our calls to action across these domains is presented in appendix exhibit 2 and discussed below.[1]

## Ensuring The Safe, Effective, And Trustworthy Use Of AI

The conceptualization, development, and deployment of AI technologies requires careful consideration to ensure their safe, effective, and trustworthy use. Trustworthiness itself is complex and encompasses concepts of fairness, equity, mitigation of bias, and sustainability. There is growing agreement that the use case needs and measurable benefits of AI technologies must first be established with a resource feasibility assessment and within the context of the societal and cultural values of the users and recipients of the technology.

Yet health AI has experienced challenges in areas such as evaluation of accuracy and reliability in settings in which tools are deployed, translation of goals into practice, problems in data management, decision errors, insufficient workflow integration, and inequitable application, among others. There are also challenges related to providing clinicians with transparency into AI analyses and the need to ensure that the operation of models is fair—that they perform equally well across patient demographics and use case–relevant subpopulations. There have been notable examples of AI performance problems, although characterizations of patient safety impacts have been very limited.[9] For example, an electronic health record vendor nationally disseminated a sepsis prediction tool that was subsequently evaluated at one hospital and was determined to detect only 183 (7 percent) of 2,552 patients with sepsis not treated with antibiotics in a timely fashion, missing 1,709 patients (67 percent) identified by clinical teams.[10] Another example is a study that showed significant racial bias in the use of a commercially developed risk calculator to allocate population health

**Democratizing access to health AI for lower-resourced organizations will promote health equity.**

resources to patients at a large academic medical center. The predictions made by the tool were biased, as they based predictions on prior inequitable health care use rather than on chronic condition–related illnesses.[11]

Demographic shifts, changes in disease incidence, and other evolving contexts thus require bias assessments for subpopulations, retraining, and tuning to incorporate new information and ensure equitable performance.[9] Predeployment testing conditions should mirror real-world deployment conditions, necessitating local validation and continuous monitoring of AI systems using "algorithmovigilance" approaches.[12,13]

For AI that qualifies as a medical device, the Food and Drug Administration (FDA) provides a rigorous framework for oversight, including postmarketing surveillance.[14,15] In other cases, stakeholders have proposed assurance laboratories to enable the multicenter evaluation of AI technologies for specific use cases.[16] There are several guidance principles, frameworks, and regulations that support AI design, development, and implementation, but there remains a strong need to harmonize, disambiguate, and address gaps among these guidance documents to ensure safe, effective, and trustworthy AI (see appendix exhibit 1).[1]

Costs associated with AI deployment pose additional challenges to equitable use and sustainability. Infrastructure and data collection require substantial capital investment. Large academic medical centers may be more able to budget for and justify the cost of relevant infrastructure compared with small or community-based health care organizations. This is a critical issue, as democratizing access to health AI for lower-resourced organizations will promote health equity. Payment mechanisms, such as new technology add-on payments from the Centers for Medicare and Medicaid Services (CMS), could provide incentives to accelerate adoption of beneficial AI, but large gaps remain in supporting the democratized AI integration into clinical care delivery.

There are also numerous challenges to fair governance of health care data, including complexities of data ownership, access rights, and sharing. It is critical to balance privacy rights with the societal need for broad data access so that AI technologies can be developed and tested using data representative of the populations in which the technologies will be used.

**PRIORITIES FOR ACTION** Although the common use of "AI" can suggest a single technology, the term actually refers to a set of technologies that can be applied in different ways and with different goals. The heterogeneity of AI technologies results in variable interpretations of what is considered AI, as interpreted from recommendations, guidelines, and regulations. Clear and concrete definitions of health care AI technologies and their applications are critical to ensuring equitable use and to providing stakeholders with a common understanding of the range of technologies, applications, and lessons learned, thereby ensuring that governance strategies are appropriately and reliably formulated.

Federal agencies should develop policies to incentivize the equitable and fair deployment of AI technologies, including policies to reduce the costs of infrastructure, data collection, and education, and to democratize access to AI deployment across a broad range of health care organizations. As leaders in health care payment innovation, CMS and other relevant agencies should consider expanding reimbursement models to encourage equitable adoption. It is also critical to require or incentivize the inclusion of patients and end users into the entire AI development and implementation life cycle.

Standardization of the AI implementation life cycle will promote transparency and ease of use of AI-enabled tools within the health care community. We suggest policies and procedures that align with best practices in disciplines relevant to the use of AI in health care: risk and change management, implementation science, human-computer interaction, and the learning health system.[17] Even when AI tools are validated and determined to have minimal bias, broad adoption requires attention to organizational integration, which poses sociotechnical barriers, including the development of new workflows or the substantial adaptation of existing ones.[18] In practice, a combination of translational science methods is necessary for effective, widespread AI use, encompassing applied informatics, process improvement, and implementation science, which focus not just on AI models and their performance but also on their impacts on care delivery. Last, reduced cost and ease of use of these frameworks is especially important for lower-resourced health systems.

Careful balance in AI governance between legislation and regulation (that is, "hard/fixed governance") and certifications and best practices (that is, "soft/fluid governance") will shape the pace and direction of AI development.[19] There is a need for a fundamentally different approach to health AI oversight that prioritizes dynamism and learning over static rules. This means moving away from rigid, predefined controls and embracing a continuous evaluation system within US regulatory frameworks that adapts to the evolving nature of AI. Some of these features are being proposed in the FDA predetermined change control framework (see appendix exhibit 1).[1] Such a system should incorporate feedback mechanisms that holistically capture both positive and negative impacts of real-world consequences. Although some static controls are necessary, such as ethical guidelines and reporting standards, application-specific approvals should be revisited as understanding of AI's risks and benefits matures.

For health AI that falls outside of evolving guidance from the FDA and the Office of the Assistant Secretary for Technology Policy (formerly the National Coordinator for Health Information Technology), we call for policy makers and regulatory agencies to support the establishment of certification processes for safe, effective, and trustworthy AI use that mature over time and provide end users with assurance of reasonable operation of AI while reducing the burden of end-user expertise in AI technologies. We suggest a standardization of governance for establishing how to define, collect, and monitor AI-related safety events in partnership with patient safety organizations.

Challenges to equity and fairness can arise on the basis of differing levels of investment in the development, validation, and fielding of valuable AI technologies. A global focus is needed across all aspects of AI design, development, and implementation on promoting equitable AI and assessing and managing bias to ensure that culturally aware health equity is realized.[20]

## Promoting The Development Of An AI-Competent Workforce

Health care personnel must be informed and discerning users of AI and active participants in establishing the value propositions and requirements of these tools. In the same way that training programs for physicians and allied health professionals require prerequisites of study in biology, chemistry, statistics, and anatomy, basic knowledge of AI and its applications is needed for all health care personnel.

Health care education underwent a transformation in recent years because of the increasing volume of and access to health care information. Health care education organizations have adapted quickly to improve learners' digital competence. Researchers and clinical groups have proposed a variety of core competencies for health care professionals' use of AI.[21,22] Beyond formal training programs, clinical and business professional societies use existing knowledge exchange platforms and dissemination pathways to train their membership on these competencies and establish new and continued education requirements and accreditation standards.

**PRIORITIES FOR ACTION** Careful integration of AI education into higher health care education will strengthen the future US health care workforce. Policy makers for higher education funding should consider incentives that support professional societies, accrediting bodies, and faculty at medical and allied health professional schools to implement new training requirements and continuous adaptation of curricula to prepare clinicians to leverage AI in patient care.

In addition, policy makers should incentivize health care educational organizations to routinely evaluate knowledge and skills to identify those that are becoming redundant as health care AI advances. Principles of parsimony must be applied to ensure that medical education programs do not not simply add training requirements for clinicians but focuses on conferring the necessary knowledge for delivering safe, effective, and compassionate care in concert with new and emerging AI technologies. For example, as AI matures, many of the technical elements and processes underpinning performance and safety may eventually be placed in the hands of certifying bodies, similar to the Clinical Laboratory Improvement Amendments, which ensured laboratory test performance as that domain matured. However, understanding of AI limitations, applications, and extensions in practice will continue to be important and should be incorporated into curricula.

In contemplating these new clinical educational requirements, however, it is essential to recognize that the US is already facing an epidemic of clinician burnout,[23] and new requirements cannot be continuously added without considering what skills are most essential for health care delivery. Although expertise in both clinical care and computer science are essential for the AI era of health care, it is unreasonable to expect every clinician to earn the equivalent of a joint degree.

We also suggest that AI education be further expanded in health-adjacent workforces, which increasingly depend on information management and digital innovation to achieve effective care delivery. Such workforces include clinical

researchers, payers, regulators, policy makers, care partners, community health workers, and behavioral coaches, among others. Investments to prepare these workers to embrace AI should include efforts to promote diversity, equity, and inclusion—for example, by reaching into new talent pools to recruit people who reflect the overall patient population.

In addition, universities could partner with graduate health care education institutions to provide health care–focused curricula to learners in ethics and equity, computer science, data science, decision analysis, and related fields. To fully capture the promise of AI in transforming health care, the US should be inclusive and promote technical experts as partners in the modern health care workforce.[24] Computer scientists, engineers, data scientists, security experts, and business administrators should be considered essential members of the health care workforce. It is important to train these learners on the clinical, regulatory compliance, and ethical principles of care delivery and information science. This should include new curricula or specialist tracks, as well as professional societies for specific disciplines supporting continuing education and upskilling.

It is imperative that all stakeholders engage in health AI to articulate and define a shared vision of what a successfully diverse and inclusive AI-competent health care workforce looks like. Only then will it be possible to implement programs to create this workforce—likely a combination of interdisciplinary generalists and highly skilled clinical, technical, and business experts, with the cultural diversity necessary to support the equitable use of AI-enabled tools.

## Promoting Research On AI In Health And Health Care

AI has emerged as a powerful tool for revolutionizing biomedical research, care delivery, and population health. Its ability to process and organize vast amounts of multiscale and multimodal data, recognize patterns therein, and make informed decisions can accelerate and improve human decision making and understanding across a broad range of problem domains.[25]

It is difficult to assess the overall US research investment in health-related AI and machine learning because many research projects use these technologies in supporting ways that are not directly captured as focused on AI. The most notable research funding agencies investing in this domain are the National Science Foundation, National Institutes of Health (NIH), Advanced Research Projects Agency for Health, Defense Advanced Research Projects Agency, and Department of Veterans Affairs, although other agencies such as the Department of Energy's National Laboratories and others have smaller portfolios. Some of the recent high-profile health-related research AI programs have been developed by the NIH and include the Artificial Intelligence/Machine Learning Consortium to Advance Health Equity and Researcher Diversity program[26] and the Bridge to Artificial Intelligence program.[27]

**PRIORITIES FOR ACTION** We call for focused research investments in the science of medicine, the practice of medicine, and the delivery of care to facilitate the safe, effective, and trustworthy use of AI across the health care sector. We summarize the likely federal administration actors relevant to policy decision making for these topics in appendix exhibit 2.[1] Consideration should be given to the crosscutting areas listed below.

Research investments in the science of medicine advance the scientific understanding of disease. There are open research questions on how AI can enable the characterization of disease mechanisms, particularly in understudied areas such as mental health, autoimmune conditions, and rare diseases. Another area of growing interest is AI's potential to accelerate drug discovery by analyzing molecular interactions, predicting drug efficacy, and identifying potential targets.

Research investments in the practice of medicine need to expand diagnosis and screening capacity, such as by exploring opportunities for AI to enable semiautomated or automated analysis of multimodal biomedical data for early disease detection. Novel approaches may enhance accuracy in diagnosing complex diseases such as cancer, cardiovascular conditions, or neurological disorders and may improve downstream outcomes. There are also unanswered questions with regard to how AI models should best be used to predict disease onset, progression, patient outcomes, or treatment responses. For example, how should AI be used in the early detection of disease if there is limited or no capacity to clinically intervene?

Research investments in the delivery of care can expand the role of AI technologies in precision medicine. Research questions remain regarding how, when, and where to leverage AI to tailor treatment based on individual patients' characteristics, genetics, and lifestyle or environmental exposures. This has broad implications for the concept of "standard of care" and how its definition or quality assessment may need to change to allow for personalized care.[28]

Another important area of research is AI integration into workflow, which requires collaboration between developers and end users.[29] It will

be important to address resistance to change, minimize workflow disruptions, and develop guidelines to evaluate the interplay of AI and workflows to ensure acceptability and adoption.[30] Last, there is a need for increased investments in research regarding health care operations, including how to employ AI to optimize workflows, resource allocation, scheduling, supply chain, and revenue-cycle operations, among others.

Balancing AI-driven decisions with human cognitive strengths and judgment also raises ethical questions. We support strong investments in the conduct of fairness, ethics, and the incorporation of cultural values into AI research, in both direct evaluation of how to achieve these goals within research and how to apply findings. Although notions of safety, effectiveness, and trustworthiness are essential in selecting tasks and contexts where AI will be used, the cultural values of the people making the choices get "embedded" in the data sets and models that are developed. Studying those effects and their downstream ramifications for a variety of stakeholders is critical.[31]

We advocate for focused research on data quality and privacy. AI relies on high-quality data. Ensuring accurate, diverse, and representative data sets is crucial, and privacy protection remains essential as well, as noted above.

It is also important to invest research in AI explainability (where appropriate). AI models often operate as "black boxes" that are difficult for human end users to interpret, and it is important to establish when, where, and how explainability is needed. However, there are settings in which the reliability of AI may be high enough and the need for explainability low enough (such as in administrative processes with appropriate performance validation, bias assessment, and ongoing surveillance) that it might not be required. Trustworthiness in the context of use underlies all principled applications of AI.

## Clarifying Responsibility And Liability In The Use Of AI

Liability for injury arising from the use of AI in medical settings is a subject of concern for physicians[31] and academics,[32] but it has received relatively little policy focus. In the United States, neither courts nor the federal government have tackled these issues directly.[33] The Bipartisan Senate AI Working Group's 2024 AI policy roadmap merely encourages relevant committees to consider the issue of liability,[34] and the 2023 AI Executive Order does not mention liability at all[35] (appendix exhibit 1).[1]

**AI is poised to transform how patients, caregivers, and health care professionals experience the management of their health, health care, and care goals.**

The liability landscape is currently muddled. It is sometimes said that AI introduces a "responsibility gap," in that those who are injured (for example, because of an adverse event from an AI-recommended treatment) or whose interests are otherwise set back (for example, when algorithmic racial bias leads to a minority group member being deprioritized) do not have a single person or entity to hold responsible. Much of medical responsibility traditionally lands on physicians, but this seems problematic when physicians cannot adequately evaluate the AI they are using and might not even be aware of how it was trained.

Legal liability has a similar problem. Existing systems parcel out liability among developers, hospitals, and clinicians in a deeply uncertain way.[36] An injured patient might try to sue any of those three entities, but the current information environment often provides insufficient evidence to know where an error originated and who can ultimately be held liable or responsible. Developers, hospitals, and clinicians can shift responsibility. This uncertainty may hinder the appropriate exercise of responsibility by key stakeholders in selecting specific AI tools to use, as well as the ability to fully realize the strengths of these tools.[37] The uncertainty may also deter some, especially in lower-resourced organizations, from adopting AI systems at all, forgoing potentially substantial benefits to patients. Courts will eventually tackle these issues, although that process takes time. In addition, legislatures may weigh in, although one-size-fits-all solutions may inadequately consider AI tool variation and care context.[38]

**PRIORITIES FOR ACTION** Policy makers should support and coordinate efforts by professional

societies to streamline the responsible adoption of medical AI by clarifying the responsibility and liability landscape for health care professionals. We propose three actions to be taken by organizations such as the National Academies, the Federation of State Medical Boards, the American Medical Informatics Association, and others: provide analyses of the most common legal questions to elucidate what clinicians and hospitals need to know and what uncertainty remains for different uses of AI, understanding that these technologies are diverse with varying liability risks; promulgate model licensing terms for medical AI that can create clearer liability rules through contract; and set model terms for indemnification or insurance against injuries involving AI. These next steps can ease the responsible adoption of AI to improve patient care.

## Conclusion

AI is poised to transform how patients, caregivers, and health care professionals experience the management of their health, health care, and care goals. Substantial challenges remain in realizing this promise. We believe that policies in four key areas can facilitate and accelerate AI in health and health care, including promoting the safe, effective, and trustworthy use of AI; promoting the development of an AI-competent health care workforce; focusing investments in key research portfolios; and clarifying AI liability and responsibilities. Through these efforts, substantial societal health improvements may be achieved. ■

The views presented in this publication are those of individual contributors and do not represent formal consensus positions of the authors' organizations, the Department of Veterans Affairs or the United States government, the National Academy of Medicine, or the National Academies of Sciences, Engineering, and Medicine. The authors thank Megan Snair, Melissa Laitner, and Emily Shambaugh for their contributions to this article. To access the authors' disclosures, click on the Details tab of the article online. [*Published online January 22, 2025.*]

## NOTES

1 To access the appendix, click on the Details tab of the article online.
2 Matheny M, Israni ST, Ahmed M, Whicher D, editors. Artificial intelligence in health care: the hope, the hype, the promise, the peril. Washington (DC): National Academies Press; 2019.
3 Rajaguru V, Han W, Kim TH, Shin J, Lee SG. LACE index to predict the high risk of 30-day readmission: a systematic review and meta-analysis. J Pers Med. 2022;12(4):545.
4 Gomes YE, Chau M, Banwell HA, Causby RS. Diagnostic accuracy of the Ottawa ankle rule to exclude fractures in acute ankle injuries in adults: a systematic review and meta-analysis. BMC Musculoskelet Disord. 2022;23(1):885.
5 Mumtaz H, Ejaz MK, Tayyab M, Vohra LI, Sapkota S, Hasan M, et al. APACHE scoring as an indicator of mortality rate in ICU patients: a cohort study. Ann Med Surg (Lond). 2023;85(3):416–21.
6 Xing Y, Sun Y, Li H, Tang M, Huang W, Zhang K, et al. $CHA_2DS_2$-VASc score as a predictor of long-term cardiac outcomes in elderly patients with or without atrial fibrillation. Clin Interv Aging. 2018;13:497–504.
7 Bayati M, Braverman M, Gillam M, Mack KM, Ruiz G, Smith MS, et al. Data-driven decisions for reducing readmissions for heart failure: general methodology and case study. PLoS One. 2014;9(10):e109264.
8 Ali O, Abdelbaki W, Shrestha A, Elbasi E, Alryalat MAA, Dwivedi YK. A systematic literature review of artificial intelligence in the healthcare sector: benefits, challenges, methodologies, and functionalities. J Innovation Knowledge. 2023;8(1):100333.
9 Finlayson SG, Subbaswamy A, Singh K, Bowers J, Kupke A, Zittrain J, et al. The clinician and dataset shift in artificial intelligence. N Engl J Med. 2021;385(3):283–6.
10 Classen DC, Longhurst C, Thomas EJ. Bending the patient safety curve: how much can AI help? NPJ Digit Med. 2023;6(1):2.
11 Obermeyer Z, Powers B, Vogeli C, Mullainathan S. Dissecting racial bias in an algorithm used to manage the health of populations. Science. 2019;366(6464):447–53.
12 Davis SE, Embí PJ, Matheny ME. Sustainable deployment of clinical prediction tools—a 360º approach to model maintenance. J Am Med Inform Assoc. 2024;31(5):1195–8.
13 Embi PJ. Algorithmovigilance—advancing methods to analyze and monitor artificial intelligence–driven health care for effectiveness and equity. JAMA Netw Open. 2021;4(4):e214622.
14 Food and Drug Administration. Software as a Medical Device (SAMD): clinical evaluation, guidance for industry and Food and Drug Administration staff [Internet]. Silver Spring (MD): FDA; 2017 Dec 8 [cited 2025 Jan 15]. Available from: https://www.fda.gov/media/100714/download
15 Food and Drug Administration. Marketing submission recommendations for a Predetermined Change Control Plan for artificial intelligence–enabled device software functions, guidance for industry and Food and Drug Administration staff [Internet]. Silver Spring (MD): FDA; 2024 Dec 4 [cited 2025 Jan 15]. Available from: www.fda.gov/media/166704/download
16 Shah NH, Halamka JD, Saria S, Pencina M, Tazbaz T, Tripathi M, et al. A nationwide network of health AI assurance laboratories. JAMA. 2024;331(3):245–9.
17 Salwei ME, Carayon P. A sociotechnical systems framework for the application of artificial intelligence in health care delivery. J Cogn Eng Decis Mak. 2022;16(4):194–206.
18 Callahan A, McElfresh D, Banda JM, Bunney G, Char D, Chen J, et al. Standing on FURM ground: a framework for evaluating fair, useful, and reliable AI models in health care systems. NEJM Catalyst [serial on the Internet]. 2024 Sep 18 [cited 2024 Dec 6]. Available from: https://catalyst.nejm.org/doi/full/10.1056/CAT.24.0131
19 Dzau VJ, Laitner MH, Temple A, Nguyen TH. Achieving the promise of artificial intelligence in health and medicine: building a foundation for the future. PNAS Nexus. 2023;2(12):pgad410.
20 Abràmoff MD, Tarver ME, Loyo-Berrios N, Trujillo S, Char D, Obermeyer Z, et al. Considerations

for addressing bias in artificial intelligence for health equity. NPJ Digit Med. 2023;6(1):170.
21 Russell RG, Lovett Novak L , Patel M, Garvey KV, Craig KJT, Jackson GP, et al. Competencies for the use of artificial intelligence–based tools by health care professionals. Acad Med. 2023;98(3):348–56.
22 Krive J, Isola M, Chang L, Patel T, Anderson M, Sreedhar R. Grounded in reality: artificial intelligence in medical education. JAMIA Open. 2023 Jun 1. [Epub ahead of print].
23 Berg S. Physician burnout rate drops below 50% for first time in 4 years [Internet]. Chicago (IL): American Medical Association; 2024 Jul 2 [cited 2024 Dec 6]. Available from: https://www.ama-assn.org/practice-management/physician-health/physician-burnout-rate-drops-below-50-first-time-4-years
24 Goldsack JC, Zanetti CA. Defining and developing the workforce needed for success in the digital era of medicine. Digit Biomark. 2020; 4(Suppl 1):136–42.
25 Haug CJ, Drazen JM. Artificial intelligence and machine learning in clinical medicine, 2023. N Engl J Med. 2023;388(13):1201–8.
26 National Institutes of Health. About the Artificial Intelligence/Machine Learning Consortium to Advance Health Equity and Researcher Diversity (AIM-AHEAD) program [Internet]. Bethesda (MD): NIH; [last reviewed 2024 Dec 12; cited 2025 Jan 15]. Available from: https://datascience.nih.gov/artificial-intelligence/aim-ahead
27 National Institutes of Health. Bridge to Artificial Intelligence (Bridge2AI) [Internet]. Bethesda (MD): NIH; [last reviewed 2024 Dec 16; cited 2025 Jan 15]. Available from: https://commonfund.nih.gov/bridge2ai
28 Lewis AE, Weiskopf N, Abrams ZB, Foraker R, Lai AM, Payne PRO, et al. Electronic health record data quality assessment and tools: a systematic review. J Am Med Inform Assoc. 2023;30(10):1730–40.
29 Jindal JA, Lungren MP, Shah NH. Ensuring useful adoption of generative artificial intelligence in healthcare. J Am Med Inform Assoc. 2024;31(6):1441–4.
30 Yu KH, Healey E, Leong TY, Kohane IS, Manrai AK. Medical artificial intelligence and human values. N Engl J Med. 2024;390(20):1895–904.
31 American Medical Association. AMA augmented intelligence research: physician sentiments around the use of AI in health care: motivations, opportunities, risks, and use cases [Internet]. Chicago (IL): AMA; 2023 Nov [cited 2024 Dec 6]. Available from: https://www.ama-assn.org/system/files/physician-ai-sentiment-report.pdf
32 Sullivan HR, Schweikart SJ. Are current tort liability doctrines adequate for addressing injury caused by AI? AMA J Ethics. 2019;21(2): E160–6.
33 Mello MM, Guha N. Understanding liability risk from using health care artificial intelligence tools. N Engl J Med. 2024;390(3):271–8.
34 Schumer C, Rounds M, Heinrich M, Young T. Driving U.S. innovation in artificial intelligence: a roadmap for artificial intelligence policy in the United States Senate [Internet]. Washington (DC): Bipartisan Senate AI Working Group; 2024 May [cited 2024 Dec 6]. Available from: https://www.schumer.senate.gov/imo/media/doc/Roadmap_Electronic1.32pm.pdf
35 Biden JR. Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence [Internet]. Washington (DC): White House; 2023 Oct 30 [cited 2024 Dec 6]. Available from: https://www.whitehouse.gov/briefing-room/presidential-actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/
36 Price WN II, Gerke S, Cohen IG. Liability for use of artificial intelligence in medicine. In: Solaiman B, Cohen IG, editors. Research handbook on health, AI, and the law. Cheltenham: Edward Elgar Publishing; 2024. p. 150–66.
37 Price WN 2nd, Gerke S, Cohen IG. Potential liability for physicians using artificial intelligence. JAMA. 2019;322(18):1765–6.
38 Price WN II, Sendak M, Balu S, Singh K. Enabling collaborative governance of medical AI. Nat Mach Intell. 2023;5(8):821–3.