AI in health care: Current applications and policy opportunities

Artificial intelligence (AI) holds the potential to improve nearly every aspect of health care by enhancing care quality, patient experience, clinical safety, affordability, and administrative efficiency. AI involves computer systems that have been trained to perform tasks that traditionally require human intelligence and judgment such as visual perception, prediction, and problem-solving.[1]

Computer scientists and data engineers have been exploring AI for over half a century; however, recent advances in hardware and software have created a more data-rich environment that has allowed AI to flourish. As more industries look to leverage AI, governments and businesses are recognizing the need to develop guidelines to safely use AI technologies and maximize benefits while addressing potential risks.

Uses in health care

AI is increasingly used across the U.S. health care ecosystem to help improve care quality. It also helps care providers navigate the growing complexity and volume of health care data. There are many different categories of AI being applied and explored across the U.S. health care system, ranging from machine learning and automated decision-making systems to natural language processing and generative AI.

AI can improve care delivery and health outcomes, including by supporting earlier disease detection, more precise treatment plans, and more effective patient engagement. AI can also streamline administrative workflows and simplify information retrieval. It can reduce workforce burden while optimizing processes like supply chain management, data quality control, and claims processing.

At Kaiser Permanente, we have developed and thoroughly evaluated AI tools like the Advance Alert Monitoring (AAM) system which saves hundreds of lives each year by using predictive analytics to identify inpatients at risk of potentially serious complications so care teams can intervene before events.[2] Our teams have designed AI systems for radiology, such as computer vision tools that can analyze mammograms and identify signs of high-risk cancers invisible to the human eye. We are also exploring AI systems for managing patient messages more efficiently connecting our members with the right care providers while monitoring emerging trends in patient requests and health needs.

How might different stakeholders use AI to improve health outcomes and care delivery?

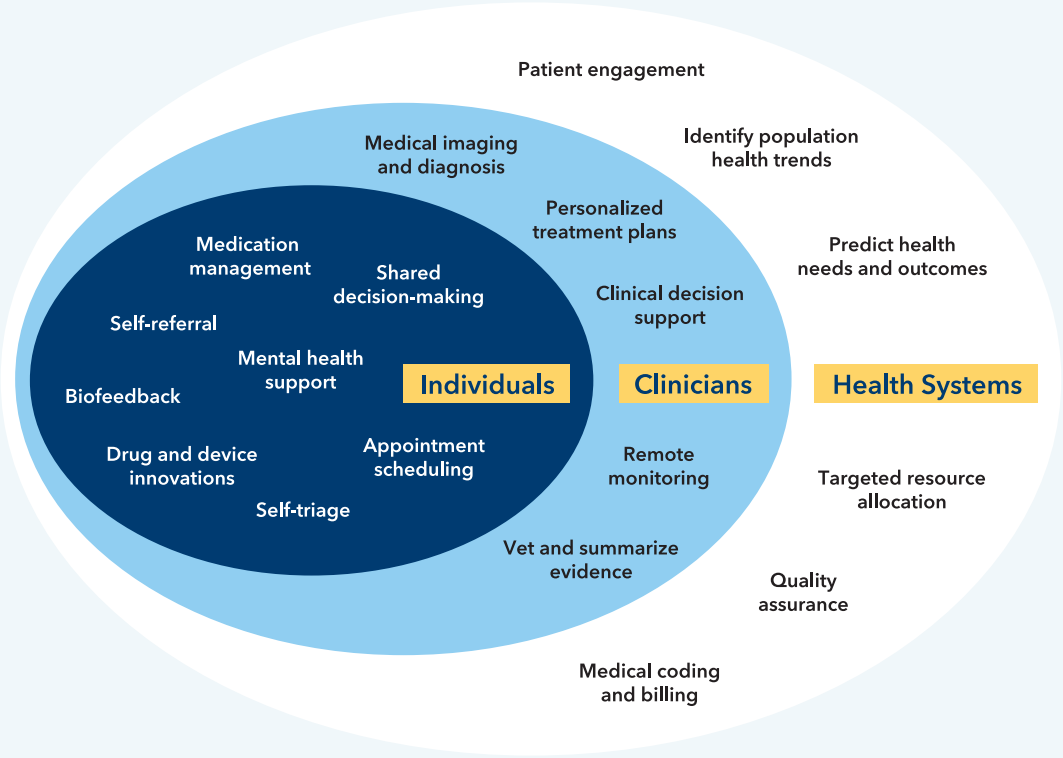

AI innovations are rapidly increasing: Trends in FDA medical device approvals

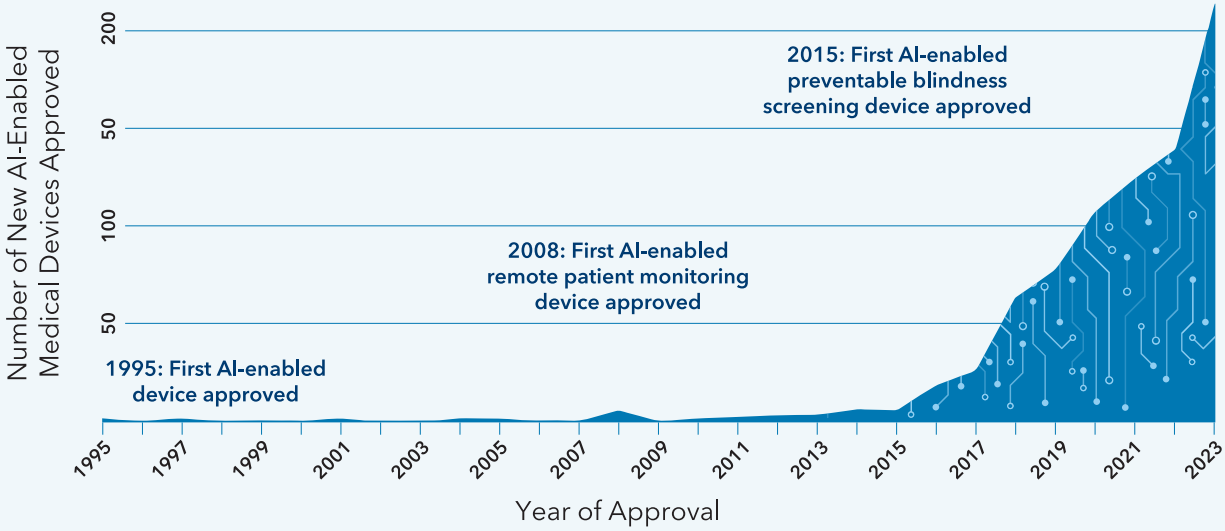

Responsible AI

As AI systems are increasingly used in health care, health care leaders must take measures to ensure the safety, reliability, equity, and trust of these tools. Generative AI models can "hallucinate," or present inaccurate or misleading information as fact,[3] and algorithms can make predictions or decisions that unintentionally but systematically disadvantage or exclude certain groups of people (known as algorithmic bias[4]). While many providers and policymakers remain hopeful about AI's potential benefits, there is also concern that AI systems could create privacy risks for patients, misdiagnose patients without sufficient oversight, and perpetuate historical inequities in health care research and delivery.

Developing and using AI systems responsibly requires that we assess their risk and value and use technology to enhance, not replace, human expertise. At Kaiser Permanente, we prioritize safety, equity, transparency, and improved health outcomes in the development, assessment, and deployment of any technology, including AI. Our AI operating model is designed to 1) Thoroughly evaluate all AI-based tools for quality, safety, and equity before use; 2) Continually assess the benefits and risks of all AI tools, regularly monitor their performance, and track improvements in patient outcomes; and 3) Safeguard patient privacy and data security.

Kaiser Permanente has a long history of developing and deploying innovative technologies at scale, so we are championing efforts to foster responsible and inclusive AI, including:

- Partnering with the National Academy of Medicine to establish an AI Code of Conduct, a guiding framework for testing, validating, and improving AI systems across health care organizations[5]
- Serving as an inaugural member of the U.S. AI Safety Institute Consortium, which is led by the U.S. Department of Commerce and aims to set standards for measuring AI safety while encouraging innovation
- Informing broadly applicable best practices for developing, using, and evaluating AI in health care as a founding member of the Coalition for Health AI, a diverse collaborative focused on equity in AI, and as a leader with the Health AI Partnership
- Working with the National Institute of Standards in Technology (NIST) to develop technical standards and a comprehensive AI risk management framework[6,7]
- Supporting the deployment and evaluation of AI across diverse health care settings and patient groups through our Augmented Intelligence in Medicine and Healthcare Initiative (AIM-HI) Center[8]
- Hosting conversations like the AI Transforming Health Care forum to encourage policymakers and health systems to collaboratively identify governance and accountability solutions

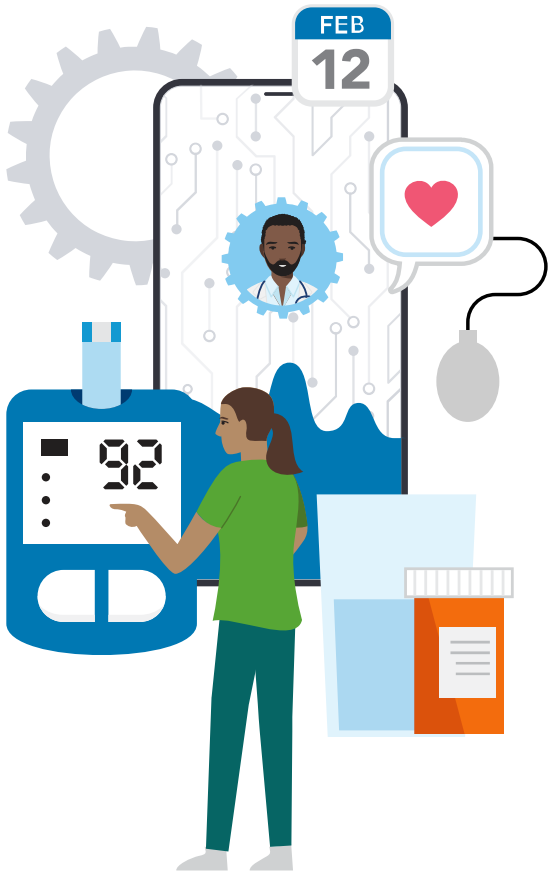

Policy opportunities

Health systems and policymakers must act to address the challenges AI presents at scale nationwide. We recommend several approaches:

Design and monitor AI systems to promote quality, safety, reliability, equity, and trust.

- **Start with equity:** Developers and health systems should incorporate equity frameworks into the AI design, evaluation, deployment, and ongoing monitoring processes. This includes using diverse, representative, and high-quality data when developing AI systems and assessing evidence of bias throughout their design and deployment.
- **Focus on outcomes:** AI performance must be specifically evaluated against safety, quality, equity, and inclusivity measures and clinical outcomes. Additionally, diverse groups of clinicians, clinical subject matter experts, and patient representatives should be included in the development and evaluation of AI systems to improve their relevance to patient experiences and key health outcomes.
- **Support rigorous research at scale:** Encourage the development of large-scale clinical trials, which would create more robust evidence to evaluate the safety and effectiveness of AI tools. This approach is also a more efficient use of resources than relying on trials done by individual health entities.
- **Evaluate long-term performance:** Quality concerns may emerge as AI systems are used with larger numbers of people, more diverse populations and settings, and over longer time periods. Long-term monitoring and follow-up are essential to understand the performance and impact of these systems after deployment. For example, establishing a post-implementation AI clinical tool surveillance program would allow health care entities to share information and best practices along with any potential issues or concerns.

- **Ensure safety and security:** Ongoing monitoring, quality controls, and safeguards are necessary to ensure the safety and security of AI systems. AI developers and health systems should work together to develop appropriate quality control systems, fail-safe systems, and feedback mechanisms.
- **Protect privacy:** Privacy must be a priority throughout the development and implementation of AI tools. Privacy protections should be consistent across the health sector and aligned with existing requirements, such as HIPAA. AI developers and health systems must address any privacy concerns that arise from the use of AI and protect personal information from unauthorized access or tampering.

Develop governance frameworks and guidelines for using AI systems responsibly.

- **Establish guidelines for appropriate use:** AI developers and health organizations should have agreed-upon use cases for the development and use of AI systems that consider the risks, benefits, limitations, and ethical considerations of each tool in a particular application. In clinical settings, humans should determine when AI is an input, subject to human review, or if AI is used to render a final analysis. Humans should always assess which conditions or caveats should apply to ensure that the AI tool is safe, effective, and used appropriately. These steps are also critical to address the risk of overly relying on AI tools across clinical and administrative settings.

- **Develop effective and equitable compliance standards:** Consider the capacity of diverse health systems, including rural hospitals and federally qualified health centers, to meet compliance expectations. Governance frameworks and standards should be appropriately rigorous but not so onerous that they prevent smaller systems with fewer resources from employing AI tools if they demonstrate they can do so safely, effectively, and equitably.
- **Promote transparency and accountability:** AI system designers and deployers should be transparent about the evidence or reasons for all outputs, as well as when, where, and how they are best used. AI tools should be sufficiently explainable while protecting the economic value of underlying algorithms and relevant intellectual property. AI systems should also enable accountability by involving human input and review across the entire design, implementation, and monitoring process.
- **Support collective action:** Policymakers and regulators should support shared benchmarking efforts and establish frameworks, working in collaboration with health systems, that can be used to evaluate the clinical implications of new AI tools and assess their readiness for implementation. Trusted health systems should lead in the identification of outcome measures and thresholds that can be used in benchmarking.

Establish national, industry-specific oversight.

- **Build in flexibility:** The health sector will benefit from a national oversight framework that is flexible and adaptable to keep pace with rapidly evolving technology. There should be industry-specific standards and guidelines for the use, development, and governance of AI systems to ensure all health care entities are measured similarly when evaluating any potential bias, error, or security issue in their AI tools.[9]
- **Emphasize function over form:** Regulatory frameworks should promote nationally recognized, technology-neutral standards, and functional, objective measures of AI inputs and outcomes. This approach balances the need for policy guardrails in the use and development of AI with fostering innovation and flexibility by avoiding regulations that dictate specific standards or technologies.
- **Coordinate standards:** Government bodies should coordinate at the federal and state levels to ensure AI standards are consistent and not duplicative or conflicting, which could stifle innovation and burden smaller developers and health systems. By working closely with health care leaders, policymakers can establish national standards that are effective, useful, and timely.

References
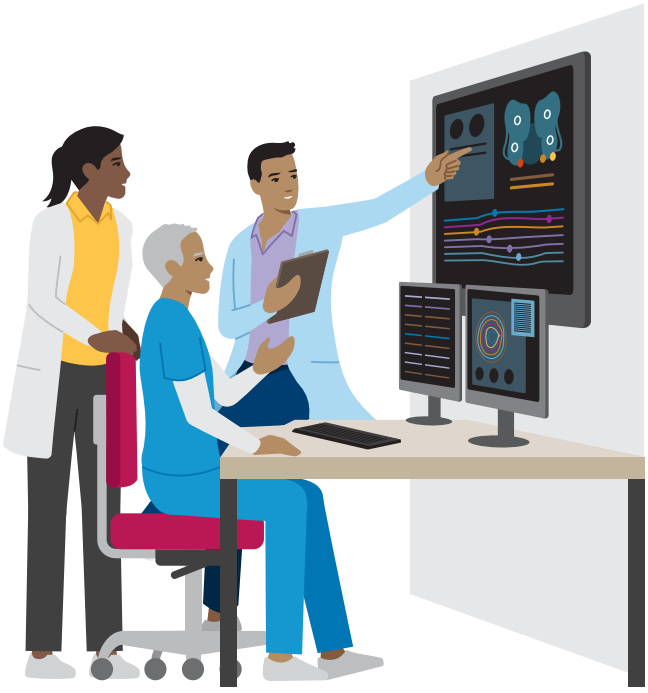
Full citations for this document can be found at: kpihp.org/references-ics.