The Federal Regulatory Landscape of AI in Health Care

Claire Tsui, Orla Levens, Lezah Calvin, & Lisa Hunter

ABOUT UNITED STATES OF CARE

United States of Care (USofCare) is a non-partisan non-profit working to ensure everyone has access to quality, affordable health care regardless of health status, social need, or income. By putting the needs of people at the forefront of our research and policy solutions, we can create a health care system that works for people.

INTRODUCTION

There is broad recognition that artificial intelligence (AI) is here to stay, and its place in health care will only expand and evolve. While AI continues to embed itself across the health care ecosystem, patients, regulators, legislators, innovators, providers and other stakeholders face an age-old challenge: nurturing innovation and advancement while mitigating the serious risks associated with unregulated technology.

This issue brief provides an overview of where AI in health care currently stands, offering background information on different types of AI, an outlook on AI utilization in health care exists today, and a snapshot of the federal regulatory landscape.

PREDICTIVE VS. GENERATIVE AI: A RUNDOWN

The term "artificial intelligence" is used in a variety of contexts, and thus has many different definitions, even across different regulatory bodies. While the use of AI in health care dates back to the 1950s, the technology has developed into increasingly more complex algorithms in the last decade, capable of making decisions and responding to the information fed into the models by users.

Predictive AI makes predictions based on the existing data fed to it and has long been used in health care going back to the early applications of the technology. Medical algorithms have been relied upon for years to assess patients for increased risk of certain conditions when given data like health history, lifestyle habits, and genetic makeup.

By contrast, **generative AI** is a newer technology that has emerged over the last decade. These models have the ability to continue learning from datasets, which trains them to improve upon interactions with patients and establish a feedback loop with health care providers. Therefore, generative AI generates new outputs (e.g. text, images, videos) based on the existing data fed to it, which it may or may not fold back into as data in the model.

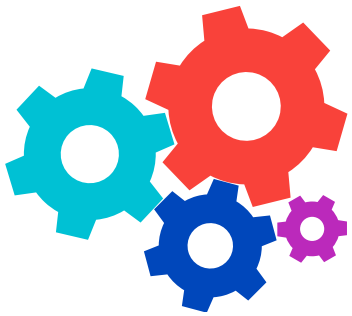

CURRENT USES OF GENERATIVE AI & INDUSTRY OUTLOOK

Given the decades-long experience the health care sector has had with predictive AI, much of the recent focus on AI in health care centers on generative AI and how it shapes health care. Put another way, generative AI is a newer frontier within the health care sphere, so it is a major area of focus for policymakers and stakeholders at this current moment.

Generative AI is mostly applied to augment two major areas of health care: provider administrative workflow and diagnostic processes. Much of what the health care industry seems to need from these algorithmic technologies isn't actually a diagnostic assistant. According to researchers at Google who developed the recent line of MedLM health care AI models, most health organizations and clinicians want AI to "help solve more back-office or logistical problems, like managing paperwork." Language processing algorithms are being used to write up comprehensive clinical summaries, freeing up time for doctors to see more patients. Chatbots help schedule appointments, manage basic patient correspondence, and conduct other telehealth functions, like outpatient monitoring.

Regardless, much of the policy and industry attention has focused on the integration of AI in diagnostic processes. **Whereas earlier AI models were historically relegated to merely analyzing and interpreting existing data, health providers can now use generative AI to take that data and propose subsequent diagnoses or treatment plans.** One of the most prominent applications so far is in imaging — in which algorithms analyze patient scans and suggest probable conditions to doctors. This technology can "detect patterns in medical images such as X-rays or MRIs" and identify diseases like cancer or Alzheimer's in their early stages. There is tentative optimism surrounding this usage of AI. Algorithms cannot get tired like human eyes do, which can potentially help "reduce diagnostic mistakes by detecting anomalies a human might overlook." More broadly, generative AIs are trained on datasets of medical records and imagery (like MRIs and CT scans) to recognize disease-related patterns. The acuity with which they perform these tasks and subsequently develop appropriate treatment proposals has been impressive. Some studies have found that even general-use AI models like ChatGPT could obtain "near-passing performance on medical exams, demonstrating a level of understanding comparable to that of a medical student."

Nonetheless, there are various shortcomings in AI that indicate that it is not ready to replace the clinical judgment and decision-making skills of human physicians. Of particular concern to patient advocates is AI's propagation of harmful, discredited race-based medicine. One particular study used interrogated four generative AI models in several rounds with nine questions designed to assess for racial bias. The results found that "all models had examples of perpetuating race-based medicine in their responses." These biases carry heavy implications for treatment decisions and patient outcomes, especially when AI tools are meant to engender a sense of objective, clinical accuracy. While this issue brief primarily focuses on generative AI, it should be noted that predictive AI and additional forms of technology have long perpetuated biases of their own. USofCare believes that future policy efforts to regulate AI must be inclusive of all forms and types of this technology, not solely generative AI. **Federal policymakers* can correct for some of these biases when regulating the safe deployment of AI health care tools.**

*There is considerable discussion about which entity or entities are responsible for correcting algorithmic biases, with some identifying developers, producers, researchers or even federal policymakers. One notable proposal is that of "assurance labs," where AI developers can test AI models according to standard criteria that would be defined with regulators, although some experts have pointed out the potential gaps of these assurance labs in addressing health equity.

Federal Activity on AI in Health Care

ADMINISTRATIVE ACTIONS

The White House released its Executive Blueprint for an AI Bill of Rights in October 2022. **So far, executive authorities have approached AI with a mixture of caution and optimism.** These algorithms and technologies, when unchecked, can be dangerous, but that danger is not inevitable. Within the Blueprint, the White House Office of Science and Technology Policy outlined five principles to guide the design, use, and deployment of automated systems to protect the American public:

Safe & Effective Systems

Automated systems will undergo "pre-deployment testing, risk identification and mitigation, and ongoing monitoring" to assess whether a system is not working as intended, possesses faulty decision-making algorithms, or explicitly seeks to violate individual safety and privacy.

Algorithmic Discrimination Protection

Automated systems risk perpetuating discriminatory practices in their judgements and decisions if there are no guardrails in place to ensure the design of more equitable algorithms. Current protective measures include "proactive equity assessments as part of the system design, use of representative data and protection against proxies for demographic features, ensuring accessibility for people with disabilities in design and development, pre-deployment and ongoing disparity testing and mitigation, and clear organizational oversight."

Data Privacy

This principle concerns automated systems' abilities to access, collect, use, transfer, and/or delete user data without knowing consent. Terms and conditions surrounding the request for data access permissions should be brief and transparent, and the current "hard-to-understand notice-and-choice practices for broad uses of data" must be reformed.

Notice & Explanation

Continuing in the same vein of AI transparency, the Notice and Explanation principle stipulates that "designers, developers, and deployers of automated systems should provide generally accessible plain language documentation" of the hows and whys AI is being implemented in a service.

Human Alternatives, Considerations, & Fallbacks

Users should be granted the opportunity to opt out of automated services in favor of human alternatives. This is particularly important in situations where "an automated system fails, it produces an error, or [the user] would like to appeal or contest its impacts on [them]." These human alternatives and fallbacks should be "accessible, equitable, effective, maintained, accompanied by appropriate operator training, and should not impose an unreasonable burden on the public."

The White House followed this Blueprint with an October 2023 Executive Order on the Safe, Secure, and Trustworthy Development & Use of Artificial Intelligence (Executive Order 14110). **This AI Executive Order called for sweeping protections for patients interacting with AI in health care systems, as well as calling for explicit guidelines to prevent and mitigate harm from algorithmic bias.**

REGULATORY ACTIONS WITHIN THE DEPARTMENT OF HEALTH & HUMAN SERVICES

The White House has put the Department of Health and Human Services (HHS) broadly in charge of overseeing the federal landscape of AI in health care. The various HHS sub-agencies are undertaking investigations within their individual domains of regulation and research, with notable efforts (as of the time of publication) including, but not limited to:

SUBAGENCY	ACTION
Office of the National Coordinator for Health Information Technology (ONC)	ONC has been designated a role in regulating the market of predictive AI in electronic health records, a space inhabited by "97% of hospitals and almost 80% of physician offices" across the US. In December 2023, ONC finalized a rule to "advance health IT interoperability and algorithm transparency."
Food & Drug Administration (FDA)	As of December 2023, the FDA "cleared, authorized, or approved more than 690 AI-enabled devices" designed to improve diagnosis and treatment processes, as well as expand patient access to care. AI also implicates changes to the existing regulatory review process for drugs, biological products, and medical devices (including software). Earlier in May 2023, the FDA also issued a discussion paper seeking guidance and feedback on navigating this new terrain. The subagency has also announced its intent to publish guidelines for data quality, transparency, and standard definitions for the use of artificial intelligence in drug and device development by the end of 2024.
National Institutes of Health (NIH)	The NIH is currently harnessing the potential of AI to research priority areas like cancer, Alzheimer's, and mental illness. Amidst this innovation, NIH has issued a Notice of Special Interest to "help improve the usability of NIH-supported data" for AI and machine learning (ML) analytics. The NIH also recently announced the Artificial Intelligence/Machine Learning Consortium to Advance Health Equity and Researcher Diversity (AIM-AHEAD) program to address the current lack of diversity in AI/ML researchers and data. These research and data gaps pose the risk of perpetuating existing medical biases and health inequities among underrepresented communities. The NIH aims to bridge these gaps by encouraging "mutually beneficial and coordinated partnerships to increase the participation and representation of researchers and communities currently underrepresented."
Office of Civil Rights (OCR)	The April 2024 final rule from the HHS Office for Civil Rights and the Centers for Medicare & Medicaid Services explicitly states that certain entities must take action to mitigate the risk of discrimination from the usage of AI tools in clinical decision making.
Agency for Healthcare Research & Quality (AHRQ)	AHRQ published a conceptual framework in December 2023 to address the applications of AI in health care, their implications in structural racial disparities in care, and the necessity of identifying and filtering out AI biases to achieve greater health equity for patients and communities. They also produced an Impact Report that same month to account for over 160 studies on the impact of health care algorithms on racial and ethnic health disparities.
Centers for Medicare & Medicaid Services (CMS)	In a similar vein to AHRQ, CMS is investigating whether the algorithms used by providers of Medicare Advantage plans to identify high-risk patients and manage costs might be introducing "inappropriate bias and restrictions" to appropriate care delivery. To address these concerns, CMS issued a Final Rule in April 2023 requiring Medicare Advantage organizations to make medical necessity determinations based on an individual's circumstances, rather than relying on algorithmic or other non-human judgment. The regulations were effective starting June 2023.
Administration for Strategic Preparedness & Response (ASPR)	ASPR is utilizing AI/ML to improve COVID-19 data collection, analysis, and forecasting, as well as the distribution of and access to vaccines.
Administration for Children & Families (ACF)	In a mix of regulation and research, ACF conducted a study focused on "emerging issues and needs" around AI in the health and human services sectors. They published the resulting report in early 2023.

A HHS Fact Sheet released in December 2023 demonstrates how major health care companies are voluntarily committing to follow the federal government's lead in navigating the potential benefits and risks of responsible AI. This includes the prioritization of advancing health equity, maintaining algorithmic transparency, expanding patient access, promoting health care affordability, improving delivery and care outcomes, and relieving workforce burnout. Furthermore, companies will establish a system of mutual accountability to ensure that their peers and partners are aligned on the AI principles established by HHS.

AGENCY RULES & REGULATIONS

In March 2024, the Office of Management and Budget (OMB) released a new policy memorandum directing "sweeping [actions] to strengthen AI safety and security, protect Americans' privacy, advance equity and civil rights, stand up for consumers and workers, promote innovation and competition." **As a core component of the landmark AI Executive Order (Executive Order 14110), the memorandum addresses how to navigate the risks of AI usage without stifling the enormous, innovative potential it stands to offer.** This approach contains five tenets:

Address Risks from the Use of AI

The memorandum mandates that federal agencies must have concrete AI safeguards by December 1, 2024. This includes methods of reliably assessing, testing, and monitoring AI's impacts on people and communities; strategies to mitigate biases being fed into algorithms; as well provisions for public transparency for how the government uses AI.

Expand Transparency of AI Use

As touched upon in the previous tenet, OMB included detailed draft instructions on how to go about public reporting.

Advance Responsible AI Innovation

OMB aims to preserve the dual interests of both caution and innovation. As of right now, adequate AI safeguards are necessary, but OMB seeks to remove unnecessary barriers to responsible innovation. OMB takes an optimistic view on AI's potential to address some of our most pressing global issues, including the climate crisis, public health, and public safety. OMB highly encourages the responsible experimentation and exploration of AI capabilities among the various federal agencies.

Grow the AI Workforce

Regulating and building safeguards requires a workforce with AI talent. The Biden Administration is committed to hiring 100 new AI professionals across federal agencies, and working alongside the Office of Personnel Management to improve talent retention within the federal government. Additionally, President Biden's Fiscal Year 2025 budget proposal contained $5 million to expand the General Services Administration's government-wide AI training program.

Strengthen AI Governance

OMB is requiring agencies to appoint Chief AI Officers and AI Governance Boards to coordinate AI usage. As of March 2024, the "Departments of Defense, Veterans Affairs, Housing and Urban Development, and State have established these governance bodies, and every CFO Act agency is required to do so by May 27, 2024."

AGENCY TASK FORCES

Executive Order 14110 created the **HHS AI Task Force**, which convenes various HHS sub-agency heads from the CMS, FDA, ONC, NIH, and CDC. Their goal is to develop a comprehensive federal strategy towards the core issues surrounding AI: drugs and devices; research and discovery; critical infrastructure; biosecurity; public health; health care and human services; internal operations; and ethics and responsibility. In doing so, HHS can set the standard of action for states and private sector actors. Additionally, many high-ranking HHS officials – including staff from CMS, ONC, OCR, CDC, and the FDA – are members of the **Coalition for Health AI (CHAI)**, a cross-stakeholder coalition convening to develop "guidelines and guardrails" for fair and transparent AI to promote high-quality health care.

CONGRESSIONAL CAUCUSES & TASK FORCES

House Speaker Mike Johnson (R-LA) and House Minority Leader Hakeem Jeffries (D-NY) appointed Representatives Jay Obernolte (R-CA) and Ted Lieu (D-CA) to lead the **Bipartisan Task Force on Artificial Intelligence.** The Task Force is composed of twelve members representing key committees of jurisdiction within the health care and technology sectors. The Task Force aims for continued AI innovation while safeguarding against new and developing threats by putting forth "guiding principles, forward-looking recommendations, and bipartisan policy proposals." Further, Representatives Troy Balderson (R-OH) and Robin Kelly (D-IL) established the House Digital Health Caucus. The mission of the Caucus is to keep policymakers informed of the rapid advancements in digital health innovation, highlight the potential impacts on patients and the health care system, and ensure that all Americans benefit from emerging digital health tools.

In the fall of 2023, the **Senate Bipartisan AI Working Group** – made up of Majority Leader Chuck Schumer (D-NY), Senator Mike Rounds (R-SD), Senator Martin Heinrich (D-NM), and Senator Todd Young (R-IN) – hosted nine insight forums with experts to develop a more comprehensive understanding of the uses and implications of AI across a variety of sectors. These forums informed the Working Group's May 2024 Roadmap for AI Policy, which outlines a framework for future legislative and executive action on AI.* While the Working Group did not explicitly endorse any legislation about the use of AI in health care, their recommendations on this topic involve using AI to promote health innovation, protect against fraud and abuse, and improve the efficiency of processes and reimbursement systems, but do not include any specific directives to center equity or the patient perspective.

What's Next?

AI will continue to embed itself throughout health care – across the industry, across health care settings, across devices, and across every community in the nation and beyond. How AI evolves as a part of the patient experience remains to be seen. While AI innovation in health care continues, so too will the federal regulatory and policy landscape. Today's AI will have a lasting impact on patients and consumers for decades. Understanding the various agencies and federal discussions taking place on AI in health care today is critical for stakeholder engagement, particularly to ensure that the perspectives of people, patients, and consumers are at the center of current and future policies.

*Notably, accountability groups released a "Shadow Report" that includes criticism in response to the Roadmap.