UC DAVIS HEALTH

manatt

MAY 2024

Developing Effective Artificial Intelligence Governance Models

Key Insights for Health Systems

Naomi Newman, Managing Director
Manatt Health
Dennis Chornenky, Chief AI Advisor
UC Davis Health
David Lubarsky, Vice Chancellor of
Human Health Sciences and CEO
UC Davis Health

Acknowledgments

With thanks to Tim Andrews, Jared Augenstein, Paul Luehr, Tom Robertson, Randi Seigel and others at Manatt and to Austin Crumpton, Principal Scientist, AI in Public Health and Medicine at Domelabs AI, for their review and expert input.

About UC Davis Health

UC Davis Health is an integrated health system with 20,000 employees, 1,000 students, 1,000 trainees, 1,300 faculty members and a large regional primary care network, providing 1.5 million-plus outpatient visits every year, with an annual budget of nearly $5 billion. UC Davis Health includes the Top 10 ranked UC Davis School of Medicine, Top 25 ranked Betty Irene Moore School of Nursing, and the Top 50 ranked UC Davis Children's Hospital and UC Davis Medical Center. The UC Davis Medical Center is a 684-bed Level 1 Trauma hospital consistently ranked among the nation's best by *U.S. News & World Report*, *Newsweek*, and *Becker's*. UC Davis Health is an innovation hub for biomedical discoveries that helps advance health around the world and is ranked in the Top 35 nationally for medical research.

About Manatt Health

Manatt Health integrates legal and consulting services to better meet the complex needs of clients across the health care system.

Combining legal excellence, firsthand experience in shaping public policy, sophisticated strategy insight and deep analytic capabilities, we provide uniquely valuable professional services to the full range of health industry players.

Our diverse team of more than 200 attorneys and consultants from Manatt, Phelps & Phillips, LLP, and its consulting subsidiary, Manatt Health Strategies, LLC, is passionate about helping our clients advance their business interests, fulfill their missions and lead health care into the future. For more information, visit https://www.manatt.com/Health or contact:

Naomi Newman
Managing Director
Manatt Health
415.291.7569
nnewman@manatt.com

Developing Effective Artificial Intelligence Governance Models

Key Insights for Health Systems

Table of Contents

Introduction

The recent explosion of artificial intelligence (AI) innovation allows us to begin to imagine a world where patients don't need a "second opinion" because the first is informed by a near-instantaneous assessment of a patient's situation based on real-time analysis of patient inputs (labs, images, notes and others) compared against the experience of millions of patients with a similar phenotype, and the latest evidence.

The application of machine learning and AI in medicine is not new, but the rapid evolution of AI technologies presents new opportunities to transform health care for the better. It also presents challenges that health systems are not yet equipped to adequately address. Today, decision-making and oversight for AI tools in health systems is fragmented—with leaders in clinical departments, IT, quality and care management, strategic sourcing, legal, compliance and finance—all grappling with whether and how to use new technologies and how to manage the risks they present. The myriad of risks AI deployment presents in the clinical setting are not trivial and, if not, addressed, may impede its rapid and effective adoption. As the state and federal level legislative and regulatory activity lag technological innovations, health systems may need to act without the benefit of a framework. Given the desire to adopt these promising and ever-evolving technologies, there is a critical need for health systems to implement robust AI governance—i.e., to establish policies and specific procedures, frameworks and ethical guidelines to ensure the responsible deployment of AI technologies within the health system. This is a need that will persist even after state and federal actors promulgate requirements and implement oversight mechanisms.

An effective AI governance model will allow health systems to take advantage of innovative technologies, while mitigating risks to patients. Health systems that learn to adapt quickly will be best positioned to incorporate these AI technologies to offer even higher quality care to patients, support a better work-life balance for clinicians and streamline costly and cumbersome administrative processes. In this paper, we provide a perspective on the steps required to step up to this challenge.

AI Is Transforming Health Care…

The launch of OpenAI's ChatGPT over a year ago unleashed a flurry of interest, product development and partnership activity in the generative AI space, including in health care. Whereas AI didn't feature prominently on the HLTH conference stages a year ago, most recently, there were at least 30 exhibitors with company names that ended in ".ai" and many main stage sessions devoted to the topic.

ChatGPT and other large language models (LLMs), including Google's Med-Gemini—specialized in medicine—can engage with humans in natural language, making the interaction seem as though there is another human speaking. The use of natural language makes information on complex topics, including health care topics, rapidly accessible, tailored and easily communicated to lay persons. Thanks to their ability to synthesize vast amounts of data, recognize patterns and generate predictions, these generative AI tools—and AI algorithms more broadly—have the power to be transformative in health care. Potential use cases are broad and multifaceted, with recent interest focusing on some of the below:

- AI scribe tools that capture interactions so medical notes can be automatically created and added to the electronic health record (EHR) after a quick review by the clinician, followed by the application of generative AI to create efficient note summaries and recommendations for the clinician. This is of high interest not only due to the economic efficiencies of automating these processes, but also because it addresses factors contributing to clinician burnout.[1]
- The use of AI algorithms to read radiology scans and flag abnormalities-helping radiologists prioritize their time on true positives and complex cases. Any pattern recognition may be better done by AI or with AI augmenting human decision-making—from recognizing rashes to counting abnormal cells on a pathology slide.
- At a more administrative level, adaptive demand algorithms can support staffing decisions and resource allocation.

While there have been a lot of attention-grabbing headlines about AI replacing physicians, the more likely scenario is that it functions as an aid—amplifying the capacity of physicians and care teams to serve more patients in more tailored ways.

Table 1: Transformative Potential of AI in Health Care Delivery & Administration

 **Health Equity**	✓ **Address Information Asymmetry** and enable greater health literacy and patient self-advocacy by providing rapid plain-language information to patients on their condition as well as access to second opinions. ✓ **Bridge Language and Cultural Barriers** by making translation tools more readily accessible at the point of care. ✓ **Improve Access to Care in Under-Served Communities:** AI chatbots, diagnostic algorithms and remote patient monitoring tools can be used to augment access to primary and specialty care in geographies that experience provider shortages.
 **Prevention**	✓ **Enhance Ability to Predict and Prevent Disease:** AI tools can analyze vast amounts of medical data, enabling more accurate predictions of disease or readmission based on risk factors that may not be readily apparent to human observers. ✓ **Improve Patient Engagement** by providing rapid access to tailored information on risk factors and strategies for patients to prevent escalation—in language that is easily accessible. ✓ **Allow for More Rapid, Data-Driven Public Health Needs Assessments and Monitoring to Guide Policy Decisions and Resource Allocation:** AI tools can be useful in predicting outbreaks and forecasting demand for public health services.

Diagnosis

- ✓ **Enhance Diagnostic Accuracy:** Misdiagnosis is a significant issue in health care—the prevalence ranges from 5% to 20%, depending on the disease or condition.[2] Misdiagnosis can have serious consequences, leading to delayed or inappropriate treatments, unnecessary procedures, increased health care costs and patient harm. Studies have shown that AI tools can exceed the diagnostic accuracy of human doctors for certain conditions. With AI, physicians have the benefit of a trained "co-pilot," and patients can benefit from timelier and data-backed diagnoses.
- ✓ **Increase Access to Specialty Diagnostics:** AI algorithms can rapidly assess diagnostic images and biomarkers for disease risk factors, allowing for greater throughput of diagnostic screenings and freeing up clinician time to focus on true positives and complex cases. People in rural, remote or otherwise under-served areas could access specialist-level review of their diagnostic tests without having to travel, thanks to AI. One health system is developing an AI-driven diagnostic tool that analyzes voice biomarkers to predict coronary artery disease.[3]

Treatment

- ✓ **Provide Efficient and Personalized Treatment Planning:** AI-driven tools have also been able to speed up the treatment planning process—for example, reducing radiation therapy treatment planning from days to mere minutes.[4]
- ✓ **Enable to Access to Bedside Decision Support:** The authors of *"Foundation models for generalist medical artificial intelligence"* describe a potential future where AI can not only provide an early warning for an adverse event, but also provide the rationale and data points to back its assessment, and recommend a course of intervention, driven by the latest clinical evidence.[5]
- ✓ **Augment Procedures:** Today, surgical robots already provide remote case proctoring. In the future, surgical robots may be able to annotate video streams of procedures in real time or raise alerts when procedure steps are omitted.[6]
- ✓ **Allow for Enhanced Remote Patient Monitoring:** Companies continue to develop the capability to monitor patients with chronic conditions remotely, through the tracking of biomarker data provided through wearable technologies or phone-based image capture. These technologies can trigger patient or clinician intervention to address concerns in real-time.
- ✓ **Provide Access to Treatment Information through Patient-Facing Chatbots:** LLMs enable the creation of chatbots that can respond to patient queries and provide guidance. There are myriad liability issues to work through, but at a basic level, robust and well-tested models can be an important source of information for patients to learn about treatment options.

Administration

- ✓ **Provide Note Transcription Tools:** AI can be used in note transcription tools to automatically convert spoken medical information into written text, potentially enhancing efficiency and accuracy in medical documentation.
- ✓ **Facilitate Demand Forecasting:** Demand forecasting tools employ predictive modeling to estimate future patient demand, enabling health care organizations to optimize staff scheduling.
- ✓ **Support Billing and Coding:** AI can be employed to automate and streamline medical billing and coding. With AI-powered automatic coding tools, medical billing professionals can spend time checking coding rather than doing the data entry manually. Revenue cycle management is a large expense for health systems.

Research

- ✓ **Identify Patients for Clinical Trials:** AI can be used to cull EHR data on patients to identify those that may be eligible for clinical trials, potentially increasing the number of historically underrepresented patients in clinical trials.
- ✓ **Support Drug Discovery and Development:** AI can be used to support drug target identification, make efficacy predictions, individualize drug response predictions and monitor drug safety and adverse events, among other things.[7]

…and Creating a Host of Potential Concerns

The rapid evolution of AI also raises many concerns, including:

- **Health Equity:** If AI algorithms are trained on biased, unrepresentative or incomplete data, they can perpetuate, or even amplify, existing biases in health care—exacerbating the risk of under-diagnosis and under-treatment for certain populations.
- **Data Leakage, Privacy and Security:** Generative AI models rely on access to large amounts of patient data for training. If not adequately protected, sensitive patient data could be vulnerable to leaks or unauthorized access, compromising confidentiality, privacy, and security of patient data.
- **Explainability:** LLMs are often described as "black boxes"—providing limited visibility into their underlying mechanics and training data. Physicians may not trust the results because they cannot validate or verify the process used to produce them.
- **Accuracy:** Generative AI can produce results that sound compelling but may be false (typically referred to as "hallucinations") and there is evidence that people are less likely to question outputs from AI relative to information from other people.[8]

- **Context Window Limitations:** The "context window" refers to how much input data an LLM can process. The most transformative applications of AI will likely involve analysis of multiple complex data sources such as EHRs, imaging systems, laboratory systems, claims systems, data from wearables other monitoring sources and daily inputs not currently considered (e.g., weather, environmental toxin exposures and heat maps of communicable diseases mapped onto personal cell phone geolocations). Current models may not have the capacity to process such large and varied data inputs and therefore may be of limited value to clinical teams.
- **Ethical:** Generative AI algorithms can produce realistic synthetic data without flagging it as such—including medical images or even patient records and claims—which could have negative effects on patient care, research or be used fraudulently.
- **Legal:** Responsibility for AI-generated outputs, intellectual property rights for training data, liability issues and regulatory compliance are complex, particularly if health systems rely on the technology for critical health care decisions.
- **Financial:** Technologies are still very much in the development and testing phase and the current hype cycle is prompting potential over-investment in unproven technologies and potential diversion of limited resources.
- **Environmental:** The energy required by AI tools and systems is ever growing and is on track to generate a carbon footprint the equivalent of some countries.

Even to the extent these risks may be mitigated, if patients and clinicians do not trust AI tools, they will be less likely to adopt them. In one recent survey, most patients expressed a desire to be informed if AI played a part in their diagnosis or treatment, *and* they felt uncomfortable receiving an AI diagnosis that could not be explained, even if it was accurate 98% of the time.[9] Professional norms play an important role in influencing the adoption of innovations in health care,[10] so if clinicians do not trust AI tools, adoption is likely to be slow. The issue of trust is complex and being better than a human who makes a mistake may be insufficient. For instance, although driverless cars are demonstrably safer overall, when there is a safety incident, there is a disproportional loss of confidence in the technology relative to when a human driver makes a similar mistake. An AI-driven clinical decision-making function may need to be 100–1,000 times better than humans to gain trust.

Current State of AI Oversight in Health Care: Multi-Faceted and Under Development

Given the rapid pace of AI advancement, and the challenges of the regulatory and legal landscape to keep pace, AI governance mandates are still largely underdeveloped. Last year, the White House published a *Blueprint for an AI Bill of Rights* intended to protect civil rights and ensure AI does not exacerbate inequities in our society.[11] In October of 2023, President Biden issued an *Executive Order on Safe, Secure and Trustworthy*

Artificial Intelligence, marking a first step in Federal oversight of AI. The Order requires developers of systems that pose a serious risk to national security, national economic security or national public health and safety to notify the government when training their models, and share their testing results with the government. It also directs the National Institute of Standards and Technology (NIST) to develop standards, tools and tests to help ensure that AI systems are safe, secure and trustworthy—several of which it just released. The U.S. Department of Health and Human Services Office of the National Coordinator for Health Information Technology finalized its "Health Data, Technology, and Interoperability" rule, which establishes a framework for the regulation of AI and other predictive algorithms incorporated into certified EHR systems used for clinical decision making. And more recently, HHS finalized antidiscrimination rules on patient care decision tools, including AI.

While there has been a flurry of activity, much of AI development in health care will likely not rise to the level of posing a threat to national security, and health care AI development is not just happening by certified EHR system vendors. To fill the void, self-regulation has emerged as a potential approach: stakeholders can self-regulate by agreeing to voluntary guardrails on safety (testing models, assessing risks, and sharing results of those assessments), security (safeguarding models against cyber threats) and trust (labeling AI-generated content)[12] until further Federal and likely State laws and regulation are enacted. Scholars from MIT's College of Computing argue for a self-regulatory organization, similar to Financial Industry Regulatory Authority (FINRA)[i] in the financial sector, for oversight of AI beyond the scope of currently regulated applications.[13]

A large group of leading academic medical centers, health systems and partner organizations have formed the Coalition on Health AI to advance the safe, effective and equitable use of AI in health care organizations. Members of this group announced a partnership with Microsoft to form a non-profit to build a nationwide network of laboratories that will test AI tools designed for use in health care.[14] And just recently, Epic launched an open-source tool for health systems to test artificial intelligence models—a tool the Health AI Partnership (a collaborative of health systems with Duke, Mayo Clinic, and Kaiser Permanente) plans to use.[15]

Developing Effective Governance Models for AI in Health Systems

While state and federal guardrails evolve and emerge over the next few years, health systems must consider how best to manage risks while tapping into the opportunities AI presents. The unique and multifaceted nature of AI means there is no single nexus of oversight within a health system leading to the risk of poorly implemented solutions, regulatory and malpractice exposures, and—at worst—patient harm. Considering these gaps, we propose steps for health systems to develop effective AI governance models in the near-term.

i. FINRA is a not-for-profit organization that works under the supervision of the Securities and Exchange Commission (an agency of the Federal Government) to write and enforce rules for brokers and broker-dealers, examine firms for compliance with those roles, foster market transparency and educate consumers—in this case investors. FINRA's Board has ten seats designated for industry members, 13 for public members and one seat for FINRA's Chief Executive Officer.

a. **Develop a prioritization process.** This can easily mimic the process to evaluate any technological or clinical decision support deployment—starting with considering the importance of the problem to be solved, impact of solving it, likelihood of a successful implementation, ease of implementation, resources required to solve the problem and maintain the solution, ROI and bandwidth of those charged with implementation (e.g., content experts, IT, administration).

b. **Bring the right experts to the table.** Managing the multifaceted risks associated with AI development and deployment requires expertise from a variety of disciplines beyond medicine, including data and computer sciences, IT, bioethics, compliance, legal and regulatory experts. Health systems should assemble an **interdisciplinary oversight committee** with representation from each of these disciplines and a patient representative.

c. **Develop an AI strategy and set of guiding principles at the enterprise level.** Create a written document that answers key questions, such as the following:

   i. Why should we use AI in the first place?

   ii. What applications and use-cases should we pursue?

   iii. How will we ensure safety, efficacy, accuracy, anti-bias security and privacy? (Including post implementation validity and monitoring)

   Broadly engaging stakeholders and those likely impacted by the adoption of AI will help produce a well-aligned strategy that provides directional guidance and authority to subsequent governance and adoption efforts. A survey by VALID AI of health systems around the country found that organizations overestimate their readiness—with many launching pilots without the strategic or governance foundation in place.[16]

d. **Take a user-centered design approach and lead with the problems, not the solutions.** Given where we are in the hype cycle, there is a risk that AI tools are implemented as the proverbial "hammer looking for a nail," resulting in leadership distraction with, and health system overinvestment in, unproven, unnecessary or poorly performing solutions. Mitigating this risk requires a clear-eyed understanding of which tasks AI is particularly well-suited to address and how they align with the organization's strategic objectives. Clinicians and staff are well-positioned to identify health system-specific opportunities for AI deployment, but in an environment of intense clinician burnout, getting them to engage with AI will require a clear value proposition for them and their patients. Often health care innovations fail to achieve widespread adoption because implementers do not engage in user-centered design with clinicians and staff and fail to consider the switching costs associated with replacing incumbent technology.[17]

e. **Consider the patient voice.** For patient-impacting AI-tools, patients should be actively engaged in considering the benefits and risks of AI deployment and in testing solutions.

f. **Inventory current use of AI-tools.** As noted, AI-enabled tools are already being deployed in health systems—and potentially without a coordinated strategy or oversight mechanism. An important early step for the oversight committee is to conduct an inventory of current AI research and deployment, and

an assessment of associated benefits and risk exposure. Compliance with HHS's new rule to strengthen non-discrimination protections and advance civil rights in health care will likely necessitate such an inventory process.

g. **Match the oversight processes and policies with the scale of the risk.** No AI tool should be deployed without human validation and confirmation—whether it be for back-office or patient-facing activities. There should be a process for model intake, registration, risk evaluation and subsequent validation and approval pathways. At intake, there should be a requirement for solution vendors to provide data cards[ii] so health system leaders can understand the sources and representativeness of training data, the modes of training used, and the extent to which humans have fine-tuned the model. Health systems should also review evidence of effectiveness of the model. The higher the risks of patient harm, the higher the degree of scrutiny of the model, and the higher the degree of testing before deployment, and of monitoring while in use.

h. **Evolve and combine emerging risk management frameworks from across health care and the technology industry.** President Biden's *Executive Order on Safe, Secure and Trustworthy Artificial Intelligence* directs the NIST to set rigorous standards for extensive red-team testing to ensure safety before public release of AI models that pose serious risks to national safety and public health. NIST has already established a Risk Management Framework organized around governing, mapping, managing and measuring risk that can provide a starting point, and recently issued four draft guidance documents.

   By breaking down the types of risks that need to be managed—from regulatory to technical to clinical and financial—health systems can leverage approaches from multiple industries to establish the frameworks and guardrails required to mitigate each type.

   - The technology policy world has been developing principles of **"Responsible AI"** to manage **regulatory, compliance and social risks**. These principles include fairness, transparency, interpretability, accountability, privacy, safety and security, human-centricity, sustainability and compliance, and continue to evolve. In health care, the sensitive nature of personal health information means that the risks of data misuse have arguably higher stakes than in many other industries and create higher reputational and compliance risk for health systems and AI tool vendors of—thereby requiring strong adherence to these principles.
   - The technology industry relies on the framework of **"ModelOps"** (short for Model Operations) to manage **technical risk**, that is, the risk the AI tool does not perform according to specifications, that the data drifts and the accuracy of the model shifts over time, or that it is not well-integrated into workflows. Model Ops refers to processes to actively manage and oversee AI models once they have been deployed within an organization, and encompasses model evaluation, monitoring and maintenance, including retraining when necessary. This also includes ensuring data security and cybersecurity.

ii. "Data Cards are structured summaries of essential facts about various aspects of ML datasets needed by stakeholders across a project's lifecycle for responsible AI development." The Data Cards Playbook - Data Cards Playbook (research.google)

- Health care has the added dimension of **clinical risk** (or what we can broadly refer to as domain-specific risk)—including erroneous diagnoses and/or treatment plans, which may result in patient harm. Managing this risk requires **clinical validation**, i.e., robust testing of solutions prior to deployment, and, importantly, on-going monitoring. It may require workflow redesign for clinicians to carefully review the recommended diagnosis and/or treatment to assess its validity.
- And there is **financial risk** where a poorly coordinated AI investment and purchasing strategy can lead to wasted resources and lost opportunities for new revenue generation. Many industries rely on a **portfolio management** framework to allow for a bird's eye view of how capital is being deployed and ensuring it is achieving results and returns for the organization.

Table 2: AI Risks and Applicable Risk Management Frameworks: Summary Table

Risk Category and Examples	Risk Management Framework
Regulatory, reputational, and compliance • Patient information is accessed or made publicly available • Use of AI model not appropriately disclosed to implicated stakeholders	**Responsible AI:** Guidance on • Fairness • Transparency • Interpretability • Accountability • Privacy • Safety and Security • Human-centricity • Sustainability • Compliance
Technical • Model does not perform according to specification • Accuracy of the model drifts over time • Model not integrate well into workflows—limiting adoption • Vulnerability of AI systems and data to manipulation	**"Model Ops"** • Model evaluation • Model monitoring • Model maintenance, including retraining when necessary • Integration • Data security • Cybersecurity
Clinical (Domain-Specific) • AI tools generate misinformation for patients • AI tools result in erroneous diagnoses, treatment plans • Disruption of clinical workflows • Misuse of data and insights that amplifies care disparities	**Clinical Validation (Domain-Specific)** • Requirements for "data cards" to ensure understanding of inclusion/exclusion criteria, source of funding, and other key determinants of the training data set • Robust testing and evaluation of solutions prior to deployment • Ongoing monitoring tools and protocols
Financial • Lack of visibility into existing AI and technology portfolio • Lack of understanding of total costs and cost impact of AI investments • Overinvestment in unproven solutions, loss of trust, and slower growth	**Portfolio Management** • Maintaining an AI inventory for visibility and management • Applying "total cost of ownership" framework • Measuring ROI to drive confidence in AI investment strategy

i. **Learn from peers.** While a few health care organizations have already made big strategic bets on AI, most are in an exploratory mode—proceeding with caution to understand potential use cases and the universe of potential solutions. UC Davis Health, together with over 50 other health systems and partner organizations around the country recently launched VALID AI, a member-led collective, to explore promising use cases, accelerate adoption of generative AI, and share experiences of implementing and scaling transformational solutions in a collaborative way.

Conclusion

Many things are true about AI in health care: It has the potential to be transformative across most, if not all, aspects of health system operations, addressing major pain points such as clinician burnout and workforce shortages, and streamlining labor intensive administrative processes. AI also has the potential to cause harm—such as perpetuating disparities in care, misguiding diagnostic and treatment decisions, and disclosing or misusing patient information. It is evolving incredibly rapidly, and many companies are developing and offering AI solutions that have yet to demonstrate wide-spread effectiveness or accuracy in clinical settings under rigorous testing conditions. Some health systems may seek to get involved early by developing AI solutions—either independently or in partnership, testing emerging vendor solutions and sharing insights with peers. Others may seek to "wait till the dust settles" to understand from the experience of peers, which solutions prove effective over time. For those seeking to be early adopters and shape the field of health care AI—standing up a multi-disciplinary group to develop guiding principles, prioritize use cases and deploy oversight frameworks—can help ensure that AI investments are well-deployed, risks are understood and managed, and lessons are quickly incorporated into a dynamic AI strategy.

1. Saag, H. S., Shah, K., Jones, S. A., Testa, P. A., Horwitz, L. I. (2019) Pajama Time: Working After Work in the Electronic Health Record. Journal of General Internal Medicine. 34(9):1695-1696. doi: 10.1007/s11606-019-05055-x.

2. Newman-Toker, D. E., Wang, Z., Zhu, Y., Nassery, N., Saber Tehrani, A. S., Schaffer, A. C., Yu-Moe, C. W., Clemens, G. D., Fanai, M., & Siegal, D. (2021), Rate of diagnostic errors and serious misdiagnosis-related harms for major vascular events, infections, and cancers: Toward a national incidence estimate using the "Big Three." Diagnosis (Berlin, Germany). https://pubmed.ncbi.nlm.nih.gov/32412440/

3. Malloy, T. (2022, April 26). *AI uses voice biomarkers to predict coronary artery disease - mayo clinic news network*. Mayo Clinic. https://newsnetwork.mayoclinic.org/discussion/ai-uses-voice-biomarkers-to-predict-coronary-artery-disease/

4. Meskó, B., Görög, M. (2020) A short guide for medical professionals in the era of artificial intelligence. *Npj Digital Medicine* **3**, 126. https://doi.org/10.1038/s41746-020-00333-z

5. Moor, M., Banerjee, O., Abad, Z. S., Krumholz, H. M., Leskovec, J., Topol, E. J., & Rajpurkar, P. (2023). Foundation models for generalist medical artificial intelligence. *National Library of Medicine, 616*(7956), 259–265. https://doi.org/10.1038/s41586-023-05881-4

6. Ibid.

7. Singh, S., Kumar, R., Payra, S., & Singh, S. K. (2023). Artificial Intelligence and machine learning in pharmacological research: Bridging the gap between Data and Drug Discovery. *Cureus*. https://doi.org/10.7759/cureus.44359

8. Spitale, G., Biller-Andorno, N., & Germani, F. (2023). Ai Model GPT-3 (dis)informs us better than humans. *Science Advances, 9*(26). https://doi.org/10.1126/sciadv.adh1850

9. Khullar, D., Casalino, L. P., Qian, Y., Lu, Y., Krumholz, H. M., & Aneja, S. (2022). Perspectives of Patients About Artificial Intelligence in Health Care. *JAMA Network Open, 5*(5). https://doi.org/10.1001/jamanetworkopen.2022.10309

10. Rebitzer, J., Rebitzer, R. (2023) Why Not Cheaper and Better? Health Care and Innovation. Oxford University Press.

11. The United States Government (2023, November 22) Blueprint for an AI Bill of Rights: Making automated systems work for the American People. https://www.whitehouse.gov/ostp/ai-bill-of-rights/

12. Shear, M. D., Kang, C., Sanger, D. E. (2023, July 21) Pressured by Biden, A.I. Companies Agree to Guardrails on New Tools, New York Times, 07-21-23. https://www.nytimes.com/2023/07/21/us/politics/ai-regulation-biden.html

13. Huttenlocher, D., Ozdaglar, A., Goldston, D. (2023) A Framework for U.S. AI Governance: Creating a Safe and Thriving AI Sector. MIT Schwartzman College of Computing. https://computing.mit.edu/wp-content/uploads/2023/11/AIPolicyBrief.pdf

14. Ross, C. (2024, January 8). Backed by Mayo Clinic and Microsoft, a nonprofit forms to test AI tools used in health care. STAT. https://www.statnews.com/2024/01/08/ai-tools-health-care-nonprofit-chai-artificial-intelligence/

15. Landi, H. (2024, May 22) Epic releases open-source AI validation tool for health systems. Fierce Healthcare. www.fiercehealthcare.com/ai-and-machine-learning/epic-releases-ai-validation-software-health-systems

16. ValidAI Newsletter (2024, January 24).

17. Rebitzer, J., Rebitzer, R. (2023) Why Not Cheaper and Better? Health Care and Innovation. Oxford University Press.

manatt