Redefining Care:
Competing Forces Shaping the Future of Medical AI

An STS Research Paper
presented to the faculty of the
School of Engineering and Applied Science
University of Virginia

by

Jeremiah Druckenmiller

March 27, 2025

On my honor as a University student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments.

Jeremiah Druckenmiller

STS Advisor: Peter Norton

Your next medical consult could be with an algorithm. According to Snowflake (2024), AI can enhance diagnostics, patient care, and operational efficiency, while saving the industry billions. Algorithms can now guide needle insertion and some surgical procedures (Knudsen, Ghaffar, & Hung, 2024). Yet AI introduces new risks, including algorithmic bias, opaque decision-making, eroded privacy, and over-reliance (Parikh, Teeple, & Navathe, 2019; Naik et al., 2022). In clinical settings, the question is not simply whether AI works, but who decides how it is used and for whose benefit.

Tech companies, hospitals, consultancies, regulators, patients, and physicians are competing to shape the standards governing medical AI. Although patients and clinicians have the most at stake in healthcare, corporations with the deepest pockets are far more influential: the very people expected to use or trust these tools will assume the greatest risk with the least say. In clinical care, diagnostics, and operational support, how AI tools are developed and deployed, and by whom, will determine whether they improve healthcare for all or entrench existing inequities.

Review of Research

Current research largely focuses on AI's inherent risks or its most promising applications, but far less attention is given to the power dynamics shaping its development and deployment. The following review provides key insights into both the potential and the pitfalls of AI in medicine—offering an essential foundation upon which to examine the major participants and their agendas.

In a commentary on machine learning in healthcare, Seneviratne, Shah, and Chu (2020) stress that addressing practical safety concerns is essential to fully harnessing AI's vast potential. They highlight two key risks: the 'black box' nature of deep learning models and the challenge

of ensuring algorithms trained on external data perform reliably in clinical settings (Seneviratne, Shah, & Chu, 2020). Success has been achieved. A 2015 study at Mount Sinai Hospital trained a deep learning model called Deep Patient on data from 700,000 individuals, enabling it to uncover hidden patterns in patient records and accurately predict diseases, including liver cancer, without human instruction. Joel Dudley, who led the team, noted that while many methods predict disease well, Deep Patient "was just way better" (Miotto et al., 2016). Similar results have been observed in cardiology where a study done at Stanford on cardiac arrhythmias found their algorithm exceeded the average cardiologist performance in both sensitivity and precision (Rajpurkar et al., 2017). A separate study on skin cancer diagnosis found that deep convolutional neural networks matched expert dermatologists in identifying both the most common and deadliest skin cancers, demonstrating AI's potential for accurate, accessible diagnostics (Esteva et al., 2017) The key uncertainty is that the algorithms provide no insight into how they achieve these results.

To peer into that black box, Google researchers modified an image recognition algorithm to reveal how it identified objects. The results, known as Deep Dream, produced surreal, alien-like visuals, and sometimes even bizarre associations—like a dumbbell always appearing with an arm affixed to it (Mordvintsev, Olah, & Tyka, 2015)—demonstrating the fundamental differences between the way humans and AI algorithms make conclusions. Clinicians are known to rely on ecological rationality, using targeted cues and expertise to make decisions with limited but relevant information. AI systems debound the clinical decision-making process, using all available information, even if it's not optimal or relevant. Debounding allows AI models to rely on irrational "shortcut" features learned from training data, mathematically improving accuracy but lacking clinical validity (Tikhomirov et al., 2024). For AI to truly assist doctors, it must

explain its predictions, as proceeding without transparency would leave a choice between blind trust and skepticism, risking patient care.

The misalignment of reasoning also makes it harder to anticipate AI errors or detect biases. Bias detection and mitigation are essential for fair and generalizable AI technology (Mittermaier, Raza, & Kvedar, 2023). These tools can amplify biases in training data, exacerbating existing inequalities in socioeconomic status, race, gender, and more, particularly disadvantaging marginalized populations with less accurate predictions or underestimated care needs (Obermeyer et al., 2019; Spector-Bagdady et al., 2022; McCradden et al., 2022). In a study by MIT and Stanford, three commercially available AI facial-analysis programs were tested on their accuracy across different skin tones and genders. They found error rates below 0.8% for light-skinned men but over 20% for darker-skinned women, reaching 34% in two of three tested algorithms. These models were tested before release, but with a dataset that was over 77% male and more than 83% white (Hardesty, 2018). A graduate student stumbled upon this bias, highlighting how unnoticed flaws could pose serious risks in healthcare.

It remains uncertain whether AI even simply as a decision aid will deliver its intended outcomes. Studies on autopilot technologies have revealed the negative effects of automated decision aids on human users. The introduction of GPS navigation and terrain displays in cockpits increased pilot distraction and cognitive load (Sarter & Schroeder, 2001; Parasuraman & Riley, 1997). Similarly, a study on computer aided detection (CAD) systems for mammography found the technology didn't enhance diagnostic accuracy and often lead to missed cancers, ultimately increasing costs for insurers without proven benefits for patients (Keane & Topol, 2018). This realization came two decades after FDA approval and widespread

use (Lehman et al., 2015). During their development, the basic question of whether these AI systems consistently and comprehensively improve patient care needs to be asked and evaluated.

Strategic Acceleration: Consultancies' Influence on for Medical AI

Major consulting firms like BCG, McKinsey, IBM, and Accenture have quickly made hundreds of millions selling AI services (Mickle, 2024). Deloitte asserts that "in this early stage of AI adoption, speed matters" and encourages companies to "implement generative AI solutions to extract immediate value, learn, and ladder-up to bolder transformation" (Levy & Lyons, 2023). Bain & Company provides similar advice, urging healthcare executives to take an "enterprise-wide" approach to AI development, warning that "otherwise, the organization can trip over itself, becoming the bottleneck to its own potential" (Berger, Sandig, & George, 2024). Deloitte recently partnered with Salesforce to develop ConvergeHEALTH, a platform they currently provide to life sciences companies. They highlight the "rapidly accelerating pace of this technology — coupled with its complexity" as a challenge for organizations, positioning themselves as the "trusted third party" needed to navigate those complexities, maximize ROI, and tailor solutions to industry needs (Dhar et al., 2023). Deloitte and Bain promote these strategies not as neutral advisors but as stakeholders with vested interests, aligning their recommendations with their consultant roles and marketing their proprietary tools to organizations without the resources to develop these technologies.

These consultancies and their beneficiary AI giants like NVIDIA, Microsoft, and Salesforce use an inevitability argument to promote medical AI (Big4 Accounting Firms, 2025; Sparks & Carringer, 2023). According to Levy and Lyons (2023), "Generative AI is an inevitability," and its efficiency and advancements are crucial, as "the cost of inaction at this

stage, particularly in a data-intensive sector like life sciences, could be significant" (Levy & Lyons, 2023). GenAI is not inherently efficient or a definitive advancement and is certainly not destined to be used everywhere; rather its benefits depend on responsible use. For example, researchers using GenAI to shortcut paper writing are producing studies with nonsensical or fabricated claims, which "jeopardize the integrity of the scientific record" and "risk undermining the basis for trust in scientific knowledge" (Haider et al. 2024). In this way, AI is dangerously degrading efficiency rather than advancing research. Current enthusiasm for AI suggests it is poised to become quite ubiquitous, with already "around 700 FDA-cleared, AI-enabled medical devices are now on the market—more than 10x the number available in 2020" (Benjamens, Dhunnoo, Meskó, 2020; FDA, 2024; Niewolny, 2024). As with many consultancies, Deloitte bends its analysis to client priorities, championing fast, lightly regulated development and quick, cost-effective implementation instead of proactive deliberation. This approach risks preventable setbacks and limits AI's potential in medicine.

AI or Obsolescence

Many participants have no choice but to hastily adopt these tools. AI is transforming drug discovery and development—a process that is notoriously slow and expensive, often taking over a decade and costing billions (Schuhmacher et al., 2023; Zhavoronkov et al., 2019). By rapidly analyzing vast stores of biological and chemical data, generative AI can identify novel therapeutic targets, design new drug molecules, and even predict the success of clinical trials. The National Human Genome Research Institute — an offshoot of the NIH — has stated, "genomics researchers need AI/ML-based computational tools that can handle, extract, and interpret the valuable information hidden" within the large troves of data now available to

scientists (NHGRI, 2022). Tools like Insilico Medicine's Pharma.AI and inClinico are already demonstrating real-world success, with AI-designed drugs advancing to phase II of their clinical trials and predictive models achieving high accuracy in their forecasting (AAAS, 2023; Nouri, 2023; Field, 2023). Here AI is increasing speed to market and making drug development for rare diseases more viable. Cheapening drug discovery has the potential to break the monopoly of legacy pharma giants, fostering innovation, increasing competition, and ultimately driving down drug prices, making life-saving treatments more accessible. To compete, big pharma must invest in its own AI systems.

Balancing Development with Privacy Concerns

Medical AI introduces opportunities for profit that may come at a cost to patients' privacy. One of AI's predicted first applications in healthcare is in administrative processes and EHR integration. "Hospitals and clinics will likely begin importing numerous AI models, …that draw on aggregate data," (Banja, 2020) in order to inform patients about a condition, schedule appointments, streamline intake processes, and compile patient records (Li 2023). These models require "vast amounts of sensitive patient data to function effectively," making them lucrative assets for companies eager to exploit and commodify this information (Cascella, 2023). Deloitte's ConvergeHEALTH product, Connect, described as a "high-touch patient engagement and support" system, is advertised to "leverage insights from patient data to demonstrate business value" (Deloitte Digital & Salesforce, 2020; Sidhu, Tobias, & Turcotte, 2020). In addition, deidentified patient data is unprotected under HIPAA, yet research shows that when combined with other data streams, "algorithms can re-identify a record from as few as 3 data points" (Crigger & Khoury, 2019). This raises ethical concerns about whether healthcare facilities should

be able to share or sell the deidentified patient data, as there are no controls on its future use. Since 2019, at least two university healthcare systems have faced lawsuits for not disclosing to patients that their personal records could be shared with or sold to private companies (MacMillan & Bensinger, 2019; Wakabayashi, 2019). This blurs the lines between healthcare support and commercial gain, and as AMA Director of Health Policy Elliott Crigger states "traditional expectations for health care privacy might no longer be attainable" (Crigger & Khoury, 2019).

Corporations have thwarted attempts at further regulation to retain ample access to patient data. In 2015, Google's AI subsidiary, DeepMind, used medical data from 1.6 million patients to develop an app that detects acute kidney injuries, raising serious transparency concerns due to the lack of patient consent. This incident led to a class-action lawsuit against Google in 2022, which was dismissed by the UK High Court in 2023, yet left unresolved questions around data transparency and patient trust (De Freitas, Rudkin, & Costello, 2024). These concerns persist as Google and other tech giants expand into the healthcare sector. Google's Chief Health Officer and ex National Coordinator for Health IT for President Barack Obama, Karen DeSalvo, stated that the company aims to transform mobile and wearable devices into continuous medical data collectors, acting as a "doctor in your pocket" (Reader, 2023). To do so the company is now proactively addressing potential regulatory challenges by hiring ex-FDA officials, such as Bakul Patel, the former chief digital health officer who shaped early AI guidelines at the FDA (Reader, 2023). They have also joined the Coalition for Health AI, "a diverse array of stakeholders [assembled] to listen, learn, and collaborate to drive the development, evaluation, and appropriate use of AI in healthcare." (CHAI, 2025; CHAI, 2023) Government officials are urgently calling for legislation to be established before AI becomes fully integrated into healthcare. Sen. Mark Warner (D-Va.) wrote in a letter to Google CEO

Sundar Pichai, "Artificial intelligence (AI) undoubtedly holds tremendous potential to improve patient care," but warned that it could "also harm patients and their data and reinforce human bias … without clear norms" (Warner, 2023). Google positions itself as a partner to the government and asserts its health division is committed to "keep[ing] your data private and secure" (Google Health, 2025). As patient data fuels more applications, the fundamental belief in the preservation of patient privacy and its societal value is challenged.

Broader Government Regulation and the Response from the Private Sector

The FDA and its Canadian and UK counterparts are emphasizing the need for regulation, advising developers to use training data representative of the intended patient populations to fight biases (FDA, 2021). According to the National Conference of State Legislatures, at least six states introduced legislation aimed at regulating AI in healthcare during the 2024 legislative session (NCSL, 2024). Colorado legislation mandates that insurers test their big data systems, including "external consumer data and information sources, algorithms, and predictive models," to ensure they do not unfairly discriminate against consumers based on race, sex, disability, or sexual orientation (Colorado DORA, 2021). A bill backed by the California Medical Association aims to ensure algorithms, AI, and other software are applied "fairly and equitably." Initially, the bill included a requirement for licensed physicians to oversee AI-based decisions to "approve, modify, or deny requests by providers," but this clause was later removed (California State Senate, 2024). Former president Biden's recent Executive Order on AI emphasizes the need for "privacy-preserving techniques" to protect personal health data and mandates "standards and best practices for detecting AI-generated content and authenticating official content," which can enhance transparency and trust in AI-driven healthcare tools. Additionally, it calls for the

Department of Health and Human Services to "establish a safety program to receive reports of—and act to remedy—harms or unsafe healthcare practices involving AI," ensuring that tools are continually monitored for safety and effectiveness (White House, 2023). While these standards are not immediately enforceable on the private sector, they can influence the broader regulatory landscape, especially as agencies begin to develop specific policies and guidelines in line with the EO. Early action by the new administration indicates the regulatory backstops may be fewer than ever. President Trump framed innovation as the priority and signaled that AI oversight should fall not to the government, but to the industry itself (Pifer, 2025).

Clear gaps remain, especially around the technologies' dynamic nature. AI is adaptive and autodidactic and as a result the FDA has considered a new approach for SaMD approvals, introducing "a total product lifecycle-based regulatory framework… that would allow for modifications to be made from real-world learning and adaptation"—but such regulation has yet to be implemented (FDA, 2017). As of 2020, the FDA did not require companies to disclose whether their technology used AI/ML when applying for approval, and as a result, some chose not to (Benjamens, Dhunnoo, Meskó, 2020). In addition, there is no federal requirement for clinicians to reveal the use of AI in patient care. HHS and CMS guidance urges AI-driven decisions be revealed as such, and states like Colorado and California have broken new ground by requiring patient disclosures or human consent in specific AI use cases, yet no FDA rule explicitly forces providers to tell patients that AI was used in their care (Silverboard, Wong, & Repetto, 2024; Canter & Brim, 2024). The AMA recently commented on the lack of a concrete regulation saying it is "still developing and will undoubtedly shift" (Smith, 2024).

Biotechnology leaders with decades of development experience are now addressing the regulatory landscape. Peter Shen, Head of Digital & Automation at Siemens Medical Solutions,

testified before the Senate, emphasizing that a “continuation of flexibility in the approval process” is crucial, warning that a “one-size-fits-all approach could seriously inhibit [AI’s] potential.” For large, financially driven institutions, a more flexible regulatory approach will accelerate adoption by reducing costs tied to lengthy approval processes. Siemens pledges to self-regulate by “openly communicat[ing] insights into underlying technology,” “carefully compil[ing] training and test datasets” for traceability, and eliminating biases, aiming to create systems that are “ethically acceptable and beneficial to humankind and society” (Shen, 2024). While this approach may be feasible for public-facing corporations, it raises concerns with smaller, less visible startups. Early-stage health-tech companies developing AI solutions continue to push boundaries and proliferate in number, with biotech venture investments nearing 2021’s record totals (Gormley, 2024).

Trade associations such as the American Hospital Association (AHA) advocate for AI tools, citing their ability to improve health outcomes with "timely and precise interventions," reducing costs, and increasing productivity at multiple stages of care. To successfully adopt AI in healthcare, the AHA advises patients and clinicians take a proactive approach: patients should engage with AI regularly, especially through tools like health chatbots, while clinicians should use AI to augment clinical decision-making (AHA, 2023). Another of these associations, AdvaMed, contends that AI should continue to be regulated as any other medical device, pointing out that the FDA’s "25 years of experience reviewing and authorizing AI/ML-enabled medical devices" has created a stable framework that supports innovation without compromising safety. Shifting AI into a separate regulatory category risks a disruption and "stifling innovation and reducing patient access" (AdvaMed, 2023). This perspective overlooks AI’s evolving nature

which likely will require ongoing and scrutinous updates to maintain relevance and accuracy—posing challenges to a one-time approval model.

Voices from the Exam Room

Doctors like Danton Char of Stanford warn that the tension between profit motives and effective healthcare in the US could lead to ethical conflicts based upon the clashing priorities of algorithm designers and clinicians (Webster, 2020; Ward, 2019). Abraham Verghese, Vice Chair for the Theory and Practice of Medicine at Stanford, cites the "greater financial incentive for relying on technology, testing, processes, and efficiency" as the driving force that is "eviscerating" the foundation of the physician-patient relationship. Physicians are "losing contact" as they engage with the "intermediary" of electronic systems. He cautions that viewing patients "through a screen" risks the omission of essential context, and subtleties that elevate the quality of care. Physicians, unlike AI, are not limited to only processing data and can perceive details—a patient's body language, tone of voice, or the worn "outline of a cigarette packet" in a pocket—revealing risks often absent from patient records (Verghese, 2016). These nuances enrich the standard of care and are exceptionally difficult for AI to interpret. Still, other clinicians see promise in AI's ability to extend care beyond clinical settings and support prevention. "Our body is giving off signals as to our health all the time… now we might be able to interpret those signals in ways that allow us to prevent disease," said Dr. Paul Friedman of the Mayo Clinic (Mims, 2018). Dr. Sanjay Aneja of Yale Cancer Center believes patients should be involved in AI adoption, stating "patient education, concerns, and comfort levels should be taken into consideration" (YSM, 2022).

Patient attitudes are nuanced, but multiple surveys and testimonials have revealed "caution remains a dominant theme" (Tyson et al., 2023). Surveyed Americans see AI as promising for healthcare but are concerned about misdiagnosis, privacy, cost, and less face time with doctors. Some patients have already benefited directly from AI-driven medical devices. For example, after years of drug- and surgery-resistant seizures, Kimberly Bari turned to a rare option—an implanted brain device powered by machine learning. The NeuroPace system, she says, "literally provides peace of mind I never imagined could exist" (Mims, 2018). Similarly, an elderly woman in New York credited an AI-powered home device, NUGU, with saving her life. After she collapsed with a brain hemorrhage, she called out "Save me… I think I'm dying," and the device automatically called an ambulance. "AI is scary… but this is a great example of how [it] can be leveraged to not only support but help and save people's lives," one lawmaker said, acknowledging both the fear and life-saving potential of the technology (Kim, 2024). Patient trust depends on the task as some uses feel acceptable, while others raise concern. In a survey done by Yale Cancer Center 54% of patients are comfortable with AI reading chest X-rays, but only 18% want AI to deliver a cancer diagnosis (Khullar et al., 2022). Another survey yielded similar results where six of ten American adults said they'd feel uncomfortable if their healthcare provider used AI for diagnosing or recommending treatments (Tyson et al., 2023). Patients nearly universally want their doctors to stay in control of AI tools and maintain a standard of informed consent (Robertson et al., 2023). There is an extreme fear of a dystopian world where "robots [are] doing surgery without any human supervision," as one advocate put it (Nichols, 2025). Notably, racial and ethnic minority groups have expressed even greater concern about misdiagnosis and privacy breaches compared to white patients (Nichols, 2025; Moy et al., 2024). Given these mixed feelings, patient advocacy groups are pushing for frameworks to protect

patient rights and build trust as AI spreads in healthcare. The Light Collective's "AI Rights for Patients" code of conduct aims to address how in the rush to adopt AI, industry efforts were "largely [conducted] in the absence of input from patient communities" (TLC, 2024; Heath, 2024). Although the Light Collective's framework echoes guidance from doctors, companies, and policymakers, there's little consensus on whether these standards are truly being met—creating a mismatch in pace as large companies race ahead while patient protections lag.

Conclusion

The debate over whether we are ready for full-scale implementation has largely been settled by action, not consensus. This signals something more than just a technological transition; it reveals how financial power, not clinical evidence or patient need, can become the deciding factor in shaping the structure and values of medical innovation. When the pace of adoption is dictated by commercial actors, the broader social contract that governs medicine—trust, consent, equity—can be rewritten without the participation of those most affected. The emergence of AI in healthcare thus offers an immensely important lesson about the fragility of public oversight in the face of private acceleration. If left unchecked, similar dynamics could apply to any field where emerging technologies collide with public welfare.

One area this paper does not explore is the use of AI in insurance, which often intersects with healthcare and may significantly impact access, cost, and fairness in care. Future research can also explore how power imbalances shape other forms of algorithmic integration, perhaps in education, employment, and public policy, and examine how regulatory or participatory frameworks might be designed to ensure innovation serves those it claims to benefit.

References

AAAS (2023, May 25). American Association for the Advancement of Science. Insilico Medicine receives IND approval for novel AI-designed USP1 inhibitor for cancer. *EurekAlert!* https://www.eurekalert.org/news-releases/990417

AdvaMed. (2023). Artificial intelligence in medtech overview. *https://www.advamed.org/member-center/resource-library/artificial-intelligence-in-medtech/*

AHA (2023). American Hospital Association. Artificial intelligence and care delivery: Market insights. https://www.aha.org/center/emerging-issues/market-insights/ai/ai-care-delivery

Banja, J. (November, 2020). How Might Artificial Intelligence Applications Impact Risk Management? *AMA Journal of Ethics*, *22*, E945-951.

Benjamens, S., Dhunnoo, P., and Meskó, B. (2020). The state of artificial intelligence-based FDA-approved medical devices and algorithms: An online database. *Npj Digital Medicine*, *3*, 118.

Berger, E., Sandig, R., and George, K. C. (2024, February). How to Successfully Scale Generative AI in Pharma. *Ban & Company*. https://www.bain.com/insights/how-to-successfully-scale-generative-ai-in-pharma/

Bhatia, S. N., and Ingber, D. E. (2014). Microfluidic organs-on-chips. *Nature Biotechnology, 32*, 760–772.

Big4 Accounting Firms (2025). Deloitte Clients. https://big4accountingfirms.com/deloitte-audit

Canter, L., and Brim, E. (2024, October 4). California Enacts Health AI Bill and Protections for Neural Data. *Covington*. https://www.insideprivacy.com/uncategorized/california-enacts-health-ai-bill-and-protections-for-neural-data/

California State Senate. (2024, March 8). Hearing on artificial intelligence in healthcare. California Senate Committee on Health. *Digital Democracy*. https://digitaldemocracy.calmatters.org

Cascella, L. M. (2023). Artificial Intelligence Risks: Data Privacy and Security. *MedPro Group*. https://www.medpro.com/artificial-intelligence-risks-privacysecurity

Ceresa, C. C., Knox, C. J., Diaz-Pinto, A., Ballotta, V., Karande, P., and Roux, B. M. (2021). Agent-based model provides insight into the mechanisms behind volumetric muscle loss injury and regeneration. *PLoS Computational Biology*, *17*.

Coalition for Health AI (2025). Our Purpose. https://chai.org/our-purpose/

Coalition for Health AI (2023, April 4). Blueprint For Trustworthy AI Implementation Guidance and Assurance for Healthcare. *MITRE Corporation*. https://chai.org/wp-content/uploads/2024/05/blueprint-for-trustworthy-ai_V1.0-2.pdf

Colorado Department of Regulatory Agencies, Division of Insurance. (2021, July 6). SB21-169: Protecting consumers from unfair discrimination in insurance practices. Colorado Department of Regulatory Agencies. https://doi.colorado.gov/for-consumers/sb21-169-protecting-consumers-from-unfair-discrimination-in-insurance-practices

Corona, B. T., Greising, S. M., Garg, K., Ward, C. L., and Walters, T. J. (2022). Retrospective characterization of a rat model of volumetric muscle loss. *BMC Musculoskeletal Disorders*, *23*, 1-13.

Crigger, E. and Khoury, C. (2019). Making Policy on Augmented Intelligence in Health Care. *AMA Journal of Ethics*, *21*(2), E188-191.

Deloitte Digital and Salesforce (2020). ConvergeHEALTH Connect™ for Life Sciences: Deliver consistent, high-touch patient engagement and support across the full treatment ecosystem. https://www2.deloitte.com/content/dam/Deloitte/us/Documents/process-and-operations/us-connect-for-life-sciences.pdf

De Freitas, I., Rudkin, T., amd Costello, E. (2024, Dec 16). Court of Appeal rejects class action claim for misuse of private information. *Farrer & Co*. https://www.farrer.co.uk/news-and-insights/court-of-appeal-rejects-class-action-claim-for-misuse-of-private-information/

Dhar, A., McCallen, M., Hittinger, T., Gosu, A., and Fera, B. (2023). From code to cure, how Generative AI can reshape the health frontier: Unlocking new levels of efficiency, effectiveness, and innovation. *Deloitte Development.* https://www2.deloitte.com/us/en/pages/life-sciences-and-health-care/articles/generative-ai-in-healthcare.html

Esteva, A., Kuprel, B., Novoa, R. A., Ko, J., Swetter, S. M., Blau, H. M., and Thrun, S. (2017). Dermatologist-level classification of skin cancer with deep neural networks. *Nature*, *542*, 115–118.

FDA (2021, Oct.) U.S. Food and Drug Administration. Good Machine Learning Practice for Medical Device Development: Guiding Principles. www.fda.gov/medical-devices/software-medical-device-samd/good-machine-learning-practice-medical-device-development-guiding-principles

FDA (2024, Dec. 20). U.S. Food and Drug Administration. Artificial Intelligence and Machine Learning (AI/ML)-Enabled Medical Devices https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-and-machine-learning-aiml-enabled-medical-devices

FDA. (2017, December 8). U.S. Food and Drug Administration Staff. Software as a Medical

Device (SAMD): Clinical Evaluation: Guidance for Industry and Food and Drug Administration Staff. https://www.fda.gov/media/100714/download

Field, H. (2023, June 29). The first fully A.I.-generated drug enters clinical trials in human patients. *CNBC*. https://www.cnbc.com/2023/06/29/ai-generated-drug-begins-clinical-trials-in-human-patients.html

Google Health (2025). Privacy Matters. https://health.google/privacy/

Gormley, B. (2024, Sept. 26). Venture Mega-Rounds Return to Biotech. *Wall Street Journal*.

Haider, J., Söderström, K. R., Ekström, B., and Rödl, M. (2024). GPT-fabricated scientific papers on Google Scholar: Key features, spread, and implications for preempting evidence manipulation. *Harvard Kennedy School Misinformation Review, 5*.

Hardesty, L. (2018, February 11). Study finds gender and skin-type bias in commercial artificial-intelligence systems. *MIT News*.

Heath, S. (2024, April 16). Patient advocates explore a patient bill of rights for AI. *Tech Target*. https://www.techtarget.com/patientengagement/news/Patient-advocates-explore-a-patient-bill-of-rights-for-AI

Kamm, R. D., and Bashir, R. (2014). Creating living cellular machines. *Proceedings of the National Academy of Sciences, 111*, 1869–1870.

Keane, P.A., Topol, E.J. (2018). With an eye to AI and autonomous diagnosis. *npj Digital Med*, 1.

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*.

Kim, C. (2024, September 21). Exclusive: AI-powered device saves Queens woman's life; how to get one for free. *Abc7 NY*. https://abc7ny.com/post/exclusive-senior-citizen-queens-says-ai-powered-device-saved-life/

Kim, H., Kwon, Y. T., Zhu, C., Wu, F., Kwon, S., Yeo, W. H., and Choo, H. J. (2021). Real-Time Functional Assay of Volumetric Muscle Loss Injured Mouse Masseter Muscles via Nanomembrane Electronics. *Advanced Science*, *8*, 21-37.

Knudsen, J. E., Ghaffar, U., Ma, R., and Hung, A. J. (2024). Clinical applications of artificial intelligence in robotic surgery. *Journal of Robotic Surgery*, *18*, 102.

Lehman, C. D. , Wellman, R. D., Buist, D. S. M., et al. (2015). Diagnostic Accuracy of Digital Screening Mammography With and Without Computer-Aided Detection. *JAMA Intern Med, 175*, 1828–1837.

Levy, V. and Lyons, P. (2023, Oct 3). Can life sciences companies unlock the full value of GenAI? *Deloitte Health Forward Blog*. https://www2.deloitte.com/us/en/blog/health-care-blog/2023/can-life-sciences-companies-unlock-the-full-value-of-genai

Li, J. (2023). Security Implications of AI Chatbots in Health Care. *Journal of Medical Internet Research*, *25*, e47551.

MacMillan, D., and Bensinger, G. (2019, November 15). Google almost made 100,000 chest X-rays public—Until it realized personal data could be exposed. *The Washington Post*. https://www.washingtonpost.com/technology/2019/11/15/google-almost-made-chest-x-rays-public-until-it-realized-personal-data-could-be-exposed/

McCradden, M. D., Anderson, J. A., A. Stephenson, E., Drysdale, E., Erdman, L., Goldenberg, A., and Zlotnik, R. (2022). A Research Ethics Framework for the Clinical Translation of Healthcare Machine Learning. *The American Journal of Bioethics*, *22*, 8–22.

Mickle, T. (2024, June 26). The A.I. Boom Has an Unlikely Early Winner: Wonky Consultants. *The New York Times*. https://www.nytimes.com/2024/06/26/technology/ai-consultants.html

Mittermaier, M., Raza, M. M., and Kvedar, J. C. (2023). Bias in AI-based models for medical applications: Challenges and mitigation strategies. *npj Digital Medicine*, *6*, 113.

Miotto, R., Li, L., Kidd, B. A., & Dudley, J. T. (2016). Deep Patient: An Unsupervised Representation to Predict the Future of Patients from the Electronic Health Records. *Scientific Reports*, *6*, 26094.

Mims, C. (2018, May 20). The AI Doctor Will See You Now. *The Wall Street Journal*. https://www.wsj.com/articles/the-ai-doctor-will-see-you-now-1526817600

Mordvintsev, A., Olah, C., and Tyka, M. (2015, June 18). Inceptionism: Going Deeper into Neural Networks. https://research.google/blog/inceptionism-going-deeper-into-neural-networks/

Morris, D. L. (2023, August 30). How technology is speeding the shift to value-based care. *Medical Economics*.

Moy, S., Irannejad, M., Jeanneret Manning, S., Farahani, M., Ahmed, Y., Gao, E., Prabhune, R., Lorenz, S., Mirza, R., & Klinger, C. (2024). Patient Perspectives on the Use of Artificial Intelligence in Health Care: A Scoping Review. *Journal of Patient-Centered Research and Reviews*, *11*, 51–62

Naik, N., Hameed, B. M. Z., Shetty, D. K., Swain, D., Shah, M., Paul, R., Aggarwal, K., Ibrahim, S., Patil, V., Smriti, K., Shetty, S., Rai, B. P., Chlosta, P., and Somani, B. K. (2022). Legal and Ethical Consideration in Artificial Intelligence in Healthcare: Who Takes Responsibility? *Frontiers in Surgery*, *9*.

NCSL (2024). National Conference of State Legislatures. Artificial intelligence 2024 legislation. *https://www.ncsl.org/technology-and-communication/artificial-intelligence-2024-legislation*

Nichols, G. (2025, January 13). How Patients Really Feel About Artificial Intelligence in Healthcare. *Leukemia & Lymphoma Society*. https://www.lls.org/blog/how-patients-really-feel-about-artificial-intelligence-healthcare

Niewolny, D. (2024, March 26). Boom in AI-Enabled Medical Devices Transforms Healthcare. *NVIDIA Healthcare News*. https://blogs.nvidia.com/blog/ai-medical-devices-gtc-2024/

NHGRI (2022, Jan 12). National Human Genome Research Institute. Artificial Intelligence, Machine Learning and Genomics. *National Institute of Health*. https://www.genome.gov/about-genomics/educational-resources/fact-sheets/artificial-intelligence-machine-learning-and-genomics

Nouri, S. (2023, Sep 5). Generative AI Drugs Are Coming. *Forbes*. https://www.forbes.com/councils/forbestechcouncil/2023/09/05/generative-ai-drugs-are-coming/

Obermeyer, Z., Powers, B., Vogeli, C., and Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, *366*, 447-453.

Parasuraman, R., and Riley, V. (1997). Humans and Automation: Use, Misuse, Disuse, Abuse. Human Factors, 39, 230-253.

Parikh, R. B., Teeple, S., and Navathe, A. S. (2019, Dec.). Addressing Bias in Artificial Intelligence in Health Care. *JAMA*, *322*, 2377–2378.

Pifer, R. (2025, March 10). Healthcare tackles AI oversight with no aid from Trump administration. *Healthcare Dive.* https://www.healthcaredive.com/news/health-artificial-intelligence-governance-trump-admin-standards/741889/

Reader, R. (2023, Dec 4). AI is driving Google's health care business. Washington doesn't know what to do about it. *Politico*. https://www.politico.com/news/2023/12/04/google-washington-ai-health-00129962

Rajpurkar, P., Hannun, A. Y., Haghpanahi, M., Bourn, C., and Ng, A. Y. (2017). Cardiologist-Level Arrhythmia Detection with Convolutional Neural Networks. *arXiv*.

Robertson C, Woods A, Bergstrand K, Findley J, Balser C, Slepian MJ. (2023, May 19). Diverse patients' attitudes towards Artificial Intelligence (AI) in diagnosis. *PLOS Digit Health*, *2*.

Sarter, N. B., and Schroeder, B. (2001). Supporting decision making and action selection under time pressure and uncertainty: the case of in-flight icing. *Human factors*, *43*(4), 573–583.

Schuhmacher, A., Hinder, M., Von Stegmann Und Stein, A., Hartl, D., and Gassmann, O. (2023, Jul 26). Analysis of pharma R&D productivity – a new perspective needed. *Drug Discovery Today*, *28*, 103726.

Seneviratne, M. G., Shah, N. H., and Chu, L. (2020). Bridging the implementation gap of machine learning in healthcare. *BMJ Innovations*, *6*, 45–47.

Sidhu A., Tobias B., and Turcotte J. (2020). Building Patient Services Programs to Deliver Better Experiences, Therapy Adherence, and Improved Health Outcomes. *Deloitte Digital*. https://www2.deloitte.com/content/dam/Deloitte/us/Documents/process-and-operations/us-building-patient-services-program.pdf

Singh, H., Meyer, A. N., and Thomas, E. J. (2018). The frequency of diagnostic errors in outpatient care: Estimations from three large observational studies involving U.S. adult populations. *BMJ Quality & Safety, 23*, 727-731.

Silverboard, D. M., Wong, M. A., and Repetto, E. B. (2024, October 31). Regulation of AI in Healthcare Utilization Management and Prior Authorization Increases. *Holland & Knight*. https://www.hklaw.com/en/insights/publications/2024/10/regulation-of-ai-in-healthcare-utilization-management

Smith, T. M. (2024, October 9). Do these 5 things to ensure AI is used ethically, safely in care. *AMA News Wire*. https://www.ama-assn.org/practice-management/digital/do-these-5-things-ensure-ai-used-ethically-safely-care

Snowflake (2024). Data Trends 2024 Healthcare and Life-Sciences. www.snowflake.com/wp-content/uploads/2024/04/Data-Trends-2024-Healthcare-and-Life-Sciences.pdf

Sparks, E., & Carringer, K. (2023, February 16). HEAVY.AI Teams with Bain, Maxar, and NVIDIA to Provide High-Fidelity Digital Twins for Telecom Networks. https://www.heavy.ai/news/heavy-ai-teams-with-bain-maxar-and-nvidia-to-provide-high fidelity-digital-twins-for-telecom-networks

Spector-Bagdady, K., Rahimzadeh, V., Jaffe, K., and Moreno, J. (2022). Promoting Ethical Deployment of Artificial Intelligence and Machine Learning in Healthcare. *The American Journal of Bioethics*, *22*, 4-7.

Tikhomirov, L., Semmler, C., McCradden, M., Searston, R., Ghassemi, M., and Oakden-Rayner, L. (2024). Medical artificial intelligence for clinicians: The lost cognitive perspective. *The Lancet Digital Health*, *6*, 589-594.

TLC (2024). The Light Collective. Collective AI Rights For Patients. https://lightcollective.org/Collective-AI-Rights-For-Patients.pdf

Tyson, A., Pasquini, G., Spencer, A., & Funk, C. (2023, February 22). 60% of Americans Would

Be Uncomfortable With Provider Relying on AI in Their Own Health Care. *Pew Research Center*. https://www.pewresearch.org/science/2023/02/22/60-of-americans-would-be-uncomfortable-with-provider-relying-on-ai-in-their-own-health-care/

Topol, E. (2019). Deep medicine: How artificial intelligence can make healthcare human again.

Verghese, A. (2016). The Importance Of Being. *Health Affairs*, *35*, 1924–1927.

Wakabayashi, D. (2019, June 26). Google and the University of Chicago Are Sued Over Data Sharing. *The New York Times*. https://www.nytimes.com/2019/06/26/technology/google-university-chicago-data-sharing-lawsuit.html

Ward, L. (2019, Oct. 14). The Ethical Dilemmas AI Poses for Health Care. *The Wall Street Journal*.

Warner, M. R. (2023, August 8). Warner Urges Google to Protect Patients and Ensure Ethical Deployment of Health Care AI. https://www.warner.senate.gov/public/index.cfm/2023/8/release-warner-urges-google-to-protect-patients-and-ensure-ethical-deployment-of-health-care-ai

Webster, P. (2020, April). Virtual health care in the era of COVID-19. *The Lancet*, *395*, 1180–1181.

The White House. (2023, October 30). FACT SHEET: President Biden issues executive order on safe, secure, and trustworthy artificial intelligence. *https://www.whitehouse.gov/briefing-room/statements-releases*

YSM (2022, May 09). ale School of Medicine. How Do Patients Feel About Using AI in Health Care? https://medicine.yale.edu/news-article/how-do-patients-feel-about-using-ai-in-healthcare/

Zhavoronkov, A., Ivanenkov, Y. A., Aliper, A., Veselov, M. S., Aladinskiy, V. A., Aladinskaya, A. V., Terentiev, V. A., Polykovskiy, D. A., Kuznetsov, M. D., Asadulaev, A., Volkov, Y., Zholus, A., Shayakhmetov, R. R., Zhebrak, A., Minaeva, L. I., Zagribelnyy, B. A., Lee, L. H., Soll, R., Madge, D., … Aspuru-Guzik, A. (2019, Sep 2). Deep learning enables rapid identification of potent DDR1 kinase inhibitors. *Nature Biotechnology*, *37*, 1038–1040.