The Competition to Determine AI's Legitimate Role in Healthcare

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

by

Sulav Regmi

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.

Sulav Regmi

STS Advisor: Peter Norton

Introduction

Artificial Intelligence (AI) is changing healthcare, offering benefits and serious challenges. While AI can improve diagnostics and patient care, it brings new risks of data security, bias, and accountability. A concern is whether AI, meant to support human judgment, ultimately displaces it, reducing the role of medical professionals.

In US healthcare, inequities in access to and quality of care are pervasive. LaViest (2023) reports that in 2018, healthcare inequities cost the U.S. economy $451 billion, up from $320 billion in 2014. Rising insurance premiums, which have increased 22% since 2018 and 47% since 2013, force many to choose between healthcare and other essential needs (2023). These financial pressures drive hospitals and insurance companies to adopt AI not solely to improve care but to cut costs.

Stakeholders—including equity advocates, hospitals, insurers, and medtech companies—are competing to shape how AI is used in healthcare. Equity advocates warn against the misuse of AI, citing cases where algorithms prematurely removed elderly patients from care facilities (Napolitano, 2021) or denied necessary medical tests in Cigna's PXDX AI system (Bendix, 2023). Healthcare providers acknowledge AI's potential, but warn of its tendency to reproduce bias (Yadav, 2023). Insurers and medtech companies push AI's "efficiency" and cost-saving capabilities, with projecting cost savings exceeding $100 billion annually (Sebastian, 2021).

The future of AI in healthcare remains uncertain. While its potential to reduce costs is undeniable, its current use raises ethical concerns. The competing agendas of equity, efficiency, and profitability will shape AI's role in healthcare for all.

Review of Research

Researchers have investigated the safety of AI in healthcare. Though large AI models require pooled data, Ross (2020) notes that healthcare data is difficult to store safely. Still, without diversity in data, systems can't be scaled to general populations. Although cloud-based storage is the best method for sharing, it creates large targets for criminal activity. Even with encryption, cloud-based storage remains vulnerable. Researchers have also raised questions about the legality of data sharing. Yadav (2023) explains that health information has different rules across jurisdictions. Sharing this data would create optimal models, but European laws regulate differently to codes such as HIPAA. Public repositories like Kaggle and The Cancer Imaging Archive offer alternatives. Publicly accessible databases, including mammographic datasets like DDSM and Optimam, enable the development of models. During the COVID-19 pandemic, these platforms helped medical AI development. Yet, as Ross notes, AI's opacity can lower trust. Li (2023) reports that many AI systems operate as "black-box models", programs that are difficult to explain even by developers. This lack of transparency influences all participants, as errors cause irreversible consequences. Li, who screened five databases of ethical journals, revealed ethical issues including violations of privacy, fairness, and transparency. Medical professionals corroborate this loss of trust. Khan (2023) contends that despite these issues, AI's use is unavoidable, as any possible benefits are too attractive to deny.

The drawbacks of AI in healthcare are not limited to safety concerns, but also stem from the poor quality of data. Yadav (2023) states that models rely on electronic health records, which often overlook groups who lack access or insurance. This underrepresentation can lead to AI suggesting incorrect actions for these communities, as it learns from data where similar patients often receive limited care. Yadav highlights similar cases, such as Amazon's recruitment

algorithm, which discriminated against women “simply because of the gender bias in the data it was trained on.” This analysis aligns with previous comments made by Yale professionals, emphasizing the incredible risk inherent in AI’s training process. Researchers and professionals throughout the medical field continue to push for stricter regulation. Jin et al. (2024) evaluated GPT-4 with Vision (GPT-4V) during medical tasks. GPT-4V achieved expert-level accuracy in a closed-book setting, outperforming physicians by nearly 4%. However, the study revealed significant shortcomings in the model’s underlying rationales. Over 20% of cases exhibited flawed image comprehension, even if the final answer was correct. This discrepancy between performance and the quality of the reasoning process reinforces earlier concerns about “black-box models”, showing a need for clarity in AI evaluations. Such findings show the broader challenges of integrating AI into clinical practice, where trust, transparency, and ethical data use remain critical.

Despite these concerns, AI has already been used heavily in treatment. Rahman et al. (2024) lists diagnosis as an application, stating AI has “greater accuracy” because of the size of its fed data. In radiology, AI systems assist in finding abnormalities in medical imaging. Rahman contends AI’s main advantage is its speed. The timeline for reporting a new foodborne illness in outbreaks is three to four weeks. An AI system could identify a set of symptoms or bacterial strains much faster, making a difference in controlling outbreaks (Rahman, 2024). Shiaelis et al. (2023) show this with a neural network that performs virus identification in less than 5 minutes.

The debate over AI’s role in healthcare is shaped by stakeholders and participants, each with an agenda. Equity advocates, hospitals, insurance companies, and medtech organizations all have a vested interest in how AI is implemented. While equity advocates push for fair and ethical access, insurers focus on efficiency and profit, creating conflict.

Equity Advocates and Healthcare Providers

Equity advocates push for fair and ethical access to healthcare. These groups cite times where AI has been misused to the detriment of vulnerable populations. This conflict is clear in cases where AI decisions override medical professionals, many resulting in legal action. In 2021, a UnitedHealthcare algorithm removed elderly patients from care facilities against doctors' opinions, creating concerns over AI's authority in patient care. The lawsuit claimed the model was “known by the company to have a 90% error rate” but was heavily used (Napolitano, 2021). Cigna faced a lawsuit for using its PXDX AI system to deny essential tests without human review. This left patients to bear unexpected costs while violating California insurance laws (Bendix, 2023). A U.S. Senate investigation found that insurers—including UnitedHealthcare, Humana, and CVS— used AI to deny requests for post-acute care. This was done at higher rates than other forms of care. Between 2019 and 2022, these insurers denied such requests at rates higher than before, Humana’s denial rate increasing by 16 times (U.S, 2024). UnitedHealthcare’s denial rate surged from 10.9% to 22.7% in these years. These increases occurred at the same time as when AI automation was used. Models include UnitedHealthcare’s “Machine Assisted Prior Authorization," designed to reduce review time, and "HCE Auto Authorization," a model that increased denials by flagging “contradicted evidence” (U.S, 2024). Both models, despite denying many cases of post-acute care, were approved by their committees.

These automated systems prioritized lowering costs over patient well-being. CVS, for example, saved $660 million in a single year by denying requests with an automated system. CVS is transparent about this, stating their previous models “jeopardized profits” (U.S, 2024). Humana admits to training its reviewers to emphasize cost when assessing care requests.

Employees were instructed to pose "surprise questions" for denial justification . This was also done “to uphold a denial on appeal” (U.S, 2024). Providing post-acute care was again labeled as high cost. PSI’s work concluded in stating “it is insurers who are using prior authorization to protect billions in profits while forcing vulnerable patients into impossible choices” (U.S, 2024). These findings show a troubling trend, as AI is being used to increase profit by denying needed medical care to patients. The investigation stated the actions taken by these companies “were primarily targeted at reducing administrative costs rather than medical costs” (U.S, 2024). Insurers externally praise AI for efficiency and aiding patient care. They argue AI allocates resources, makes less mistakes, and approves treatments. This investigation shows a different motive, one where AI is used to cut costs at the patient's expense. This puts insurance executives directly at odds with patients who rely on timely and necessary medical care. As denials increase, so do legal challenges as patients and advocates push back against AI systems.

In another case, the Texas Attorney General (AG) settled with AI firm Pieces Technologies. The firm was found to have made deceptive values about the accuracy of its AI tools, ones used in hospitals. The group had claimed their product had a "severe hallucination rate" of less than 1 per 100,000, misrepresenting high reliability (Fox, 2024). However, the AG's investigation found these metrics were likely inaccurate and may have misled hospitals regarding the safety of the AI products. The AG stated that “AI companies offering products used in high-risk settings owe it to the public and to their clients to be transparent about their risks, limitations and appropriate use” (Fox, 2024). This statement highlights the risks of deploying AI in healthcare without transparency. In the case of Pieces Technologies, the company pushed an extremely low error rate, yet their tools were found to be potentially unreliable. Despite this, the system had already been adopted in four hospitals, where it handled sensitive patient data and

influenced care decisions (Fox, 2024). This case shows how quickly flawed AI can enter clinical settings without proper oversight. When these tools are inaccurate or misleading, they risk deepening healthcare disparities. Without regulation and honest disclosure, AI can lead to harmful outcomes under the guise of innovation. This ongoing struggle shows the conflicting agendas shaping AI's role in healthcare - one driven by corporate interest, the other by need of ethical treatment.

Medical Professionals and Hospitals

Medical professionals are also core participants in this debate. Yale experts warn that AI systems can inherit and amplify biases, misrepresenting and harming minority groups (Backman, 2023). This is a shared concern between researchers and practitioners, stemming from training data. Professionals note that without strong oversight, AI could promote inequities rather than resolve them. Many push that AI tools are "already harming minoritized communities." Biased algorithms have been identified "that require racial or ethnic minorities to be considerably more ill than their white counterparts to receive the same diagnosis, treatment, or resources" (Backman, 2023). These algorithms are pervasive, present in multiple specialties, including cardiac surgery and kidney transplantation. Insurers counter these arguments and defend their use of AI systems. The Reinsurance Group of America is one such insurer, pushing AI's potential to save over $100 billion annually through automation. They contend that this benefits both patients and providers (Sebastian, 2021). These groups claim AI hastens administrative work, streamlines claims processing, and reduces human error in decisions.

Hospitals keep an outwardly neutral stance. The American Hospital Association accepts AI's advantages but calls for oversight to ensure responsible use (2024). However, AI can reduce

operational costs for hospitals, causing incentive to use such systems. By using AI, hospitals aim to be competitive in a changing healthcare industry. CitiBank claims automation could cut 25%-30% of the administrative costs. These costs make up a quarter of U.S. healthcare spending (Healthcare Admin Costs, n.d). As a loaning agency, CitiBank focuses on financial efficiency, and hospital stakeholders see value in these claims as they struggle with rising costs. These claims clash with real-world experience. Abraham Verghese (2018) noted that new professionals are "shocked to find that the focus on the ward doesn't revolve around the patients but around the computers." He described doctors as "the highest-paid clerical worker in the hospital." While investors praise AI's efficiency, its impact on doctors is clear — more patient turnover, less autonomy, and rising burnout. This account shows that AI in healthcare isn't just saving money. AI often negatively changes how doctors work, causing a disconnect between patient and professional. The claims of CitiBank or insurers ignore this, pushing financial data with no context.

The American Medical Association screened physicians and found they were concerned with data privacy and poor integration with current systems (AMA, 2025). Professionals are vocal about the risks of AI in healthcare, more so when used without adequate testing or clinician input. In April 2024, over 200 nurses from the California Nurses Association protested outside Kaiser Permanente's San Francisco Medical Center, expressing concerns over the use of untested AI systems. Michelle Gutierrez Vo, a registered nurse and CNA president, stated, "But what we are witnessing in our hospitals is the degradation and devaluation of our nursing practice through the use of these untested technologies." Cathy Kennedy, another CNA president, emphasized, "No patient should be a guinea pig and no nurse should be replaced by a robot" (Blum, 2024). Issues raised include AI chatbots potentially misdirecting patients during

emergencies, AI tools issuing false alarms, and staffing algorithms failing to account for essential tasks. Nurses like Micheal Kennedy report that hospital leadership, backed by tech investors, “always want to be on the cutting edge of everything. And so when something new comes out, they’re going to jump right on it. I think that’s part of why they dive headfirst into this AI thing” (Kennedy & Cockerell, 2024). These dynamics have eroded autonomy and created burnout, as clinicians feel they are “operators of the machines” rather than skilled practitioners. Kennedy (2024) states that the given reason for these tools “is always that it will make life easier for us,” but contends that hospitals are “squeezing more out of us, so they can ultimately hire fewer of us.” AI is again presented as a universal solution, but negatively affects groups when added to existing systems.

Despite these statements, AMA (2025) states that physicians have “less apprehension” about AI systems than they did before. Providers see the potential in AI, especially with transparency. The AMA emphasizes that AI should assist, not replace, human judgment in clinical settings.

Future Technologies

Medtech firms are deeply invested in AI’s expansion. GitHub’s Copilot X and Mintlify make generative AI to accelerate software engineering. Paige and Pictor Labs claim their AI tools enhance diagnostic precision (Schroer, 2024). Accenture members push for AI wearable items, wanting a "mass-market makeover" in chronic disease management. In particular, they target diabetes and cardiovascular conditions (Kawalec, 2024). These firms view AI as more than a tool, using it as a market disruptor. These products are made for the benefit of their users, but also embed AI into more product sectors, creating large profits.

Generative AI is projected to grow faster in healthcare than any other industry, likely to reach a $22 billion market by 2027. With a compound annual growth rate of 85%, continuous investment is made to generate profit (Schroer, 2024). Boston Consulting Group (BCG) identified over 60 use cases across R&D, operations, sales and marketing, customer support, and corporate functions ranked by impact and ease of implementation. BCG warns that delaying use risks long-term disadvantage and urges rapid pilot launches under GenAI frameworks. BCG pushes their motive to “empower the organization and individuals to jointly benefit from generative AI” (Schroer, 2024). KPMG (2024) forecasts the generative AI healthcare market to top $21.7 billion by 2032, again making a case regarding monetary gains. Deloitte (2020) and MedTech Europe estimate that eight AI categories—wearables, imaging, and laboratory applications being a few—could save 400,000 lives, deliver €200 billion in annual savings, and free 1.8 billion healthcare hours each year. These groups contend GenAI’s addition to healthcare is “ the equivalent of having 500,000 additional full time health care professionals” (Deloitte, 2020). These groups also mention improving patient outcomes, but focus their resources on predicting savings and finding markets to approach.

The spread of medtech firms is more effective when supported. Despite ongoing debates, public opinion leans optimistic about AI’s role in healthcare. A national survey of 926 subjects shows that over half of American respondents believe AI will improve healthcare in the next five years (Khullar, 2022). Transparency is a key factor shaping public acceptance. The survey indicates that 66% of respondents consider it "very important" to be informed when AI plays a big role in their diagnosis or treatment, with an additional 29.8% viewing it as “somewhat important.” This suggests that while the public welcomes AI advancements, they also demand clear communication about its influence on healthcare decisions. This, as seen in CVS or

UnitedHealthcare, is not being done. Life changing decisions are being made through automation without knowledge of the patient it concerns.

Comfort levels with AI-generated diagnoses depend on accuracy and explainability. When an AI system correctly diagnoses 90% of cases but cannot explain its reasoning, 40.5% of respondents feel uncomfortable, while only 28.5% feel comfortable. However, when accuracy increases to 98%, discomfort drops to 39.5%, and comfort rises to 42% (Khullar, 2022). These values represent public opinion, a large factor that influences the future of participant groups in healthcare. Public opinion shows that confidence in AI depends on performance and accountability. Medtech firms developing AI systems have an opportunity to align their work with these expectations. By focusing on accuracy and transparency these firms can use public optimism to drive AI adoption in healthcare.

Progress is shown in how hardware and simulation-based design are being used to ensure safety, accuracy, and better care. NVIDIA's (2025) *Isaac for Healthcare* is a "physical AI platform" that enables robotic systems to learn in "a physically accurate virtual environment" designed for surgery and imaging. NVIDIA is adamant that testing and validation can be achieved virtually "before deployment in the physical world." This project shows how AI, when carefully designed and tested, can truly extend care. Products like this prioritize caution over fast inclusion into already safe systems, a stark difference from AI's current use.

Conclusion

AI is reshaping healthcare, bringing both promise and risk. While it can promote speed, reduce costs, and enhance diagnostics, its flaws raise serious ethical concerns. Insurers use AI to cut expenses, often at the cost of patient care. Equity advocates warn that these systems can reinforce long-standing biases, leading to the denial of necessary treatment. Researchers and

medical professionals stress the need for oversight. Groups note that AI tools are trained on historical datasets that may be incomplete, biased, or unrepresentative of diverse patient groups. Hospitals and medtech firms see AI as an opportunity for growth. They see potential in using AI to optimize workflows, reduce provider burden, and personalize patient care. However, concerns are brought to impacts on provider patient relationships. AI tools often push their role of supplementing human judgement, some models overriding the actions of professionals. The push for automation favors cutting costs, but the human cost—patient trust, medical autonomy, and care quality cannot be ignored.

Despite concerns, AI is rapidly being integrated into healthcare. As calls for regulation and transparency grow louder, some providers are gaining confidence in AI. The future of AI in healthcare will depend on how these competing economic, ethical, and clinical interests shape its use. Without safeguards, the pursuit of efficiency may come at too great a cost.

References

AHA (2024, May). American Hospital Association. AHA urges need for flexibility in regulation of AI in health care. *AHA News.*

AMA (2025, Feb 12). Physician enthusiasm grows for health care AI. *American Medical Association.* www.ama-assn.org/press-center/press-releases/ama-physician-enthusiasm-grows-health-care-ai

Backman, I. (2023, Dec 22). Eliminating racial bias in health care AI: Expert panel offers guidelines. Yale School of Medicine. medicine.yale.edu/news-article/eliminating-racial-bias-in-health-care-ai-expert-panel-offers-guidelines

Bendix, J. (2023, Aug 7). Cigna using AI to reject claims, lawsuit charges. *MedicalEconomics*. www.medicaleconomics.com/view/cigna-using-ai-to-reject-claims-lawsuit-charges

Blum, K. (2024, Aug 9). California nurses protest 'untested' AI as it proliferates in health care. *AHCJ*. healthjournalism.org/blog/2024/08/california-nurses-protest-untested-ai-as-it-proliferates-in-health-care/

Deloitte. (2020, Oct). The socio-economic impact of AI in healthcare.

Fox, A. (2024, September 20). Texas AG settles with clinical genAI company. *HealthcareITNews*. www.healthcareitnews.com/news/texas-ag-settles-lawsuit-clinical-genai-company

Healthcare Admin Costs: How AI Can Ease the Pain. (n.d.). www.citigroup.com/global/insights/healthcare-admin-costs-how-ai-can-ease-the-pain-

Jin, Q., Chen, F., Zhou, Y. et al. Hidden flaws behind expert-level accuracy of multimodal GPT-4 vision in medicine. *npj Digit. Med.* 7, 190 (2024)

Kawalec, T., Richards, O., & Karaca-Griffin, S. (2024, May 29). Reinventing MedTech with intelligent technologies. *Accenture*. www.accenture.com/us-en/insights/life-sciences/reinventing-medtech-intelligent-technologies

Kennedy, M., Cockerell, I. (2024, Nov 12). I'm a neurology ICU nurse The creep of AI in our hospitals terrifies me. *.coda* www.codastory.com/surveillance-and-control/nursing-ai-hospitals-robots-capture/

KFF (2023, Dec. 8). 2023 Employer Health Benefits Survey

Khan, B., Fatima, H., Qureshi, A., Kumar, S., Hanan, A., Hussain, J., & Abdullah, S. (2023). Drawbacks of Artificial Intelligence and Their Potential Solutions in the Healthcare Sector. *Biomedical Materials & Devices (New York, N.Y.),* 1-8. Advance online publication.

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), e2210309.

KPMG. (Sep 2024). Klynveld Peat Marwick Goerdeler. MedTech at the crossroads: Challenges and opportunities.

Kulikowski, C. A. (2015). An Opening Chapter of the First Generation of Artificial Intelligence in Medicine: The First Rutgers AIM Workshop, June 1975. *Yearbook of Medical Informatics, 10(1),* 227-233.

LaVeist, T. A., Pérez-Stable, E. J., Richard, P., Anderson, A., Isaac, A., Santiago, R., Okoh, C., Breen, N., Farhat, T., Assenov, A., & Gaskin, D. J. (2023). The economic burden of racial, ethnic, and educational health inequities in the US. *JAMA, 329(19),* 1682.

Li, F., Ruijs, N., & Lu, Y. (2023). Ethics & AI: A systematic review on ethical concerns and related strategies for designing with AI in healthcare. *AI, 4(1),* 28-53.

Napolitano, E. (2023, Nov 21). UnitedHealth uses faulty AI to deny elderly patients medically necessary coverage, lawsuit claims. *CBS News.* www.cbsnews.com/news/unitedhealth-lawsuit-ai-deny-claims-medicare-advantage-health-insurance-denials/

NVIDIA (2025, Mar 18). NVIDIA and GE HealthCare Collaborate to Advance the Development of Autonomous Diagnostic Imaging With Physical AI (press release). nvidianews.nvidia.com/news/nvidia-and-ge-healthcare-collaborate-to-advance-the-development-of-autonomous-diagnostic-imaging-with-physical-ai

Poalelungi, D. G., Musat, C. L., Fulga, A., Neagu, M., Neagu, A. I., Piraianu, A. I., & Fulga, I. (2023). Advancing Patient Care: How Artificial Intelligence Is Transforming Healthcare. *Journal of Personalized Medicine, 13(8),* 1214.

Rahman, M. A., Victoros, E., Ernest, J., Davis, R., Shanjana, Y., & Islam, M. R. (2024). Impact of Artificial Intelligence (AI) Technology in Healthcare Sector: A Critical Evaluation of Both Sides of the Coin. Clinical pathology (Thousand Oaks, Ventura County, Calif.), 17, 2632010X241226887.

Ross, P., & Spates, K. (2020). Considering the Safety and Quality of Artificial Intelligence in Health Care. *Joint Commission Journal on Quality and Patient Safety, 46(10),* 596-599.

Schroer, D., Simon, S., & Trommer, G. (2024, May 14). Medtech's Generative AI opportunity. *BCG Global.* www.bcg.com/publications/2023/generative-ai-in-medtech

Sebastian, D. (2021, Sep). Artificial Intelligence and health insurance. www.rgare.com/knowledge-center/article/ai-and-health-insurance

Shiaelis, N., Tometzki, A., Peto, L., McMahon, A., Hepp, C., Bickerton, E., Favard, C., Muriaux, D., Andersson, M., Oakley, S., Vaughan, A., Matthews, P. C., Stoesser, N., Crook, D. W., Kapanidis, A. N., & Robb, N. C. (2023). Virus Detection and Identification in Minutes Using Single-Particle Imaging and Deep Learning. *ACS Nano*, *17*(1), 697–710.

U.S. Senate Permanent Subcommittee on Investigations. (2024, Oct 17). Refusal of recovery: How Medicare Advantage plans deny post-acute care.

Verghese, A. (2018, May 28). How tech can turn doctors into clerical workers. *The New York Times.* www.nytimes.com/interactive/2018/05/16/magazine/health-issue-what-we-lose-with-data-driven-medicine.html

Yadav, N., Pandey, S., Gupta, A., Dudani, P., Gupta, S., & Rangarajan, K. (2023). Data Privacy in Healthcare: In the Era of Artificial Intelligence. *Indian Dermatology Online Journal, 14(6),* 788-792.