Overreliance on Algorithms:
Competition over Medical AI in Clinical Care

A Sociotechnical Research Paper
presented to the faculty of the
School of Engineering and Applied Science
University of Virginia

by

Eric Thomas

April 9, 2021

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.

Eric Thomas

Sociotechnical advisor: Peter Norton, Department of Engineering and Society

Overreliance on Algorithms: Competition over Medical AI in Clinical Care

"There's nothing that I've seen in my 30-plus years studying medicine that could be as impactful and transformative," says physician Topol, referring to the current integration of AI in global healthcare (Szabo, 2019). As such, physicians, hospitals, insurers, patient advocacies, and med-tech vendors compete to draw the line between legitimate and excessive reliance on medical AI in clinical care. Physician Nissen worries that "it's only a matter of time before something like this leads to a serious health problem," which has already come close with an AI incorrectly predicting that pneumonia patients are less likely to die if they have asthma (Matheny et al., 2019; Szabo, 2019). As technology, such as electronic health records, has disrupted healthcare, physician Verghese remarks how "the real patient in the bed often feels neglected, a mere placeholder for the virtual record," which is a consequence of rushing to adopt technology without properly applying it (Verghese, 2011). While physicians are hopeful that clinical care will be improved by AI that is being integrated by hospitals, insurers, and med-tech vendors, physicians, patients, and patient advocacies are pushing for a better balance between human and computer involvement.

Review of Research

Though Tonekaboni et al. (2019) and Watson et al. (2019) highlight the value of explainable medical AI, they do not fully examine the counterpoint that explainability may not be needed if there is overwhelming proof of the AI's efficacy and that the physician relegates its prediction as extra evidence for their explanation and decision. While Angehrn et al. (2020) discuss AI applications and how poor data quality compounds bias, explainability is not examined. Though London (2019) analyzes the pros and cons of accuracy versus explainability, other AI issues, such as patient privacy, are not discussed. The aforementioned research is mainly from the perspective of health and AI experts, lacking analysis of how patient viewpoints affect the problem. Arnold (2021) thoroughly reviews the ethical implications of AI, but concerns about major tech vendors entering healthcare and the resulting consequences, such as data collection, are exempt.

Discussion by Topol & Verghese (2020) have motivated the preservation of social rituals among healthcare professionals and patients and showed the potential of AI improving that if AI is applied properly. Through careful thought by Wang et al. (2020) of what use cases where explainability may be needed, the idea of keeping AI options open is stressed.

Social Rituals Between Healthcare Professionals and Patients

Some physicians emphasize that AI should be applied to suitable, time-consuming tasks such that physicians can devote more time to key social rituals with other healthcare workers and patients. Disappointed with how electronic health records are competing with the patient for the doctor's attention, physician Verghese observes how "rituals are about transformation . . . the cementing of the doctor-patient relationship—a way of saying: 'I will see you through this illness. I will be with you through thick and thin,'" which is "key to earning their trust" (Verghese, 2011). While AI lacks this human dimension, Verghese notes that AI can free "more time for human-to-patient interaction" by automating tasks, such as medical charting, which can "improve care and allow physicians to record, and accurately register, more phenotypes and more nuance" (Verghese et al., 2018). Physician Ofri says that "we need to have the time" to care for and examine the patient, "but time that is not filled in by other tasks that are easily piled on," referring to tasks that AI excels at, such as memorizing every type of disease (Topol et al., 2019). While physicians Arnold and Komesaroff admit that "if patients come to reinterpret technologies as 'healers, initiates, or prophets,' then it is distinctly possible that traditional therapeutic rituals will cease to operate," they suggest that "if physicians are able to remain in control of these devices and techniques then they will be able to deploy them as transactional symbols" (Arnold et al., 2020). Physician Topol claims that "the greatest gift that AI can give us is to go 'back to the future' —to get us to the humanity in medicine, which is presence, a physical exam, listening, and building trust, communication, and deep empathy" (Topol & Verghese, 2020).

Responsibility to Explain AI's Predictions

To avoid excessive use of AI, some physicians and patient advocacies assert that AI should be required to explain its predictions. According to physician Thingalaya (2019), "we need an Algorithm Transparency initiative as the first step for us to engage in the evidence-based use of health care AI tools," implying that transparency is vital in proving the AI's benefits. Physician Eckert says that "the explanation associated with this prediction can also be helpful as the clinician makes the decision, so with transparency the clinician can actually see and understand the factors associated with the prediction" (Teredesai et al., 2018). Since "physicians will seek insights into what AI is doing and won't accept a black box when it comes to patient safety, quality and outcomes," the American Hospital Association declare that "transparency in the AI algorithm is necessary to understand what went into it and why" (AHA, 2019). Though Alliance for AI concede that there is a trade-off between explainability and accuracy, they say that AI should be explainable "when fairness is critical" (e.g., cancer screenings), "when consequences are far-reaching" (e.g., risky operations), and "when transparency is required by law" (Allgood et al., 2019). Physicians Kaushal and Khullar think that "it seems sensible to demand levels of explainability from different types of AI," referring to cases where explainability is valued, such as "abstract or unexplored problems, in which likelihood of bias is high, or when such models [AI] are the primary drivers of diagnostic and treatment decisions" (Wang et al., 2020). Physician Sgaier promotes casual AI that explains cause and effect, arguing that "causal AI identifies the underlying web of causes of a behavior or event and furnishes critical insights that predictive models fail to provide" (Sgaier et al., 2020).

Rather than require AI to explain its predictions, some physicians stress that physicians themselves should interpret and explain the AI's predictions, provided that the AI has been proven to improve outcomes. According to physicians Verghese et al. (2018), "clinicians should seek a partnership in which the machine predicts (at a demonstrably higher accuracy), and the human explains and decides on action," disagreeing with the notion that "models that are not causal and cannot explain the underlying process are useless." Remarking how "transparency in and of itself may not be necessary," physician Char says that, for example, "we don't know how a therapy works . . . but we can demonstrate that it reliably produces the desired therapeutic effect and that the effect can be monitored," implying that whether the AI works and improves outcomes should be prioritized over knowing how it works (Ward, 2019). Comparing AI to drug regulation, physician Obermeyer says that "explainability is not really part of the FDA's criteria," clarifying that "we just care that it works" (Ward, 2019). According to London (2019), "modern clinicians prescribed aspirin as an analgesic for nearly a century without understanding the mechanism through which it works," likening AI to medicine. Because "discrete and known tasks" have a "low probability of incorporating unseen social biases . . . where human clinicians can readily intervene," physicians Kaushal and Khullar state that they "may be relatively comfortable with using black box models for image analysis, laboratory testing, or natural language processing of clinical notes" (Wang et al., 2020). Physician Kohane says that "we should be more worried about what is a false positive rate and what is a false negative rate over time of these programs, so that even if they are used in black box fashion like the rest of medicine, they are sufficiently reliable" (Bender, 2019). Though physicians may not understand AI predictions, physician Lehman suggests that, if AI is proven to work, then "that itself can teach us new explanations" and "we just have to think differently" (Bender, 2019).

In response to criticism, some med-tech vendors are publicizing their support and providing tools for explainable AI. Since "it is crucial for an organization to have a full understanding of the AI decision-making processes with model monitoring and accountability of AI and not to trust them blindly," IBM says that "explainable AI also helps promote end user trust, model auditability and productive use of AI" and "mitigates compliance, legal, security and reputational risks of production AI" (IBM, n.d.). As one of Google's "objectives for AI," Google affirms that it "will design AI systems that provide appropriate opportunities for feedback, relevant explanations, and appeal" and are "subject to appropriate human direction and control" (Google, n.d.a). Since "the human may feel that the model is wrong, when in fact it is right" and that "we [humans] don't always reason rationally," Google claims that explainability is important for "empowering the human" and helping them understand AI predictions (Google, n.d.b). To promote explainable AI, Microsoft provides InterpretML, a tool that allows developers to "train interpretable glassbox models and explain blackbox systems" (Microsoft, 2021).

Patients' Perspectives of Medical AI

Much of the public holds AI in a positive light, showing that they as patients are interested in AI when it is applied to suitable use cases. For general health care, 54 percent of people state that they are "willing to engage with AI and robotics for their healthcare needs" (PwC, 2016). About 80 percent of patients say that they are "ready for use of AI without human control" or willing to "accept AI only with human control" when presented with four use cases: skin cancer analysis, flare-prevention monitoring, smart clothes for physical therapy, and symptom assessment (Tran et al., 2019). Additionally, patients are 55 to 66 percent likely to use AI for certain use cases, such as blood testing, virtual coaching, DNA analysis for genetic health risks, and virtual nursing (Accenture, 2018). In skeletal radiography, patients rated their confidence in AI-assisted interpretation and management at 7.06 and 4.86 out of 10, respectively, suggesting that patients are assured about AI diagnoses (York et al., 2020). Some patients assert that AI is a "useful tool," is "able to catch things humans can't," and would be a "step forward" because they "will make things better" by "improv[ing] speed and quality of data analysis," while one patient remarked how "we've been waiting for a faster identification of things and this can only help" (Dousa, 2020).

While some patients have positive views of AI, other patients still prefer recommendations from a human physician instead of a computer, believing AI to be inferior and less personalized. Since patients say that AI care is less unique, they hesitate to use services that are provided by a computer even when the computer has a higher accuracy rate than the human physician (Longoni et al., 2019). When deciding whether to undergo an operation, patients say that they are less likely to follow AI recommendations and feel less responsible for the decision when it is recommended by a physician (Promberger & Baron, 2006). Patients note that they trust both recommendations and decisions more when they are from a human physician

(Promberger & Baron, 2006). Though 75 percent of surveyed patients said that they trust physicians, only 31 percent trust diagnoses or treatments given entirely by AI (Accenture, 2020). Patients cite reasons, such as “I like visiting the doctor” and “AI might not understand me properly,” for not using AI services, which suggests that they think human care is more aware of their needs (Accenture, 2018). Some patients agreed that since AI represents “the prioritization of the collective above the individual,” it is questionable whether “there will remain a place in society for individual vulnerabilities” (Laï et al., 2020). One patient thinks that “it is comforting to have a human around you” and “to have a human be the bridge between robotics and the person,” concluding with “I think personally it’s still nice to get some human element of care” (Dousa, 2020). Another patient asserts that “it should be the way it’s done now, they give you all the options and the patient can make the decisions,” clarifying that the decision-maker should not be “the machine or anyone else” (Dousa, 2020).

Patient Privacy

As major tech vendors enter healthcare with AI services, some patients and patient advocacies express concern over the resulting data collection. Some of the public refer to “immaturity of AI technology” and “distrust of related companies” as two chief reasons why they dislike medical AI, believing that “AI doctors might invade personal privacy and lack legal supervision” (Gao et al., 2020). Some patients agree that the “notion of ‘non-belonging health data’ appeared to be problematic,” suggesting that patient privacy laws are unprepared for AI (Laï et al., 2020). 72 percent of surveyed people say that they are willing to share their health data with their physician, 53 percent with their health insurer, 12 percent with a government organization, and 11 percent with a tech company (DeSilva et al., 2020). Only 31 percent of people claim that they trust Apple with their health data, whereas other tech companies fared worse (Wronski, 2020). In addition to many patients asserting that more regulatory safeguards are needed to protect them from corporate surveillance and targeted advertising, especially with entities beyond their healthcare provider and insurer, one patient says that “before I share my data, I would really need to interrogate the company and its aims” (Dousa, 2020). The American Hospital Association argues that “patient privacy is a key concern and affects how AI is developed and tested,” so “individual privacy must be protected at all times” (AHA, 2019).

Some physicians are skeptical of AI that is offered by major med-tech vendors, implying that they would favor stronger patient privacy and AI alternatives that are developed only by hospitals. Physician Pearl argues that no services by major med-tech vendors “will make a substantial impact where it matters most: on the quality and cost of U.S. healthcare,” because “no major tech company is willing to accept medical liability,” and “tech companies will face major data-ownership issues ahead” (Pearl, 2019). When asked if med-tech vendors can be trusted with patients’ data, physician Topol says that “HIPAA isn’t enough to protect our data,” and “if the person doesn’t own their data and get it as inputs, then artificial intelligence is markedly diminished in terms of benefits for that person” (Farr, 2019). After noting how “Google’s DeepMind . . . caused controversy by acquiring data from 1.6 million NHS patients in a partnership with a London hospital without patients’ personal consent,” physician Elwyn asserts that “sensitive clinical data must be protected from vulnerability to hacking, inappropriate distribution, and microtargeting: ownership, storage, and sharing systems will need to be much better defined, described, and regulated” (Elwyn et al., 2018). Since “the digitized clinical encounter . . . interrogated by AI systems will arrive soon,” physician Elwyn proposes that “it is time to develop policies on how to collect, store, manage, and share these resources and to maximise the value of the data to patients and to many other stakeholders" (Elwyn et al., 2018). Furthermore, surveyed physicians preferred AI that would be offered by a hospital system or a government institution over a tech company (Landi, 2019).

Though major med-tech vendors and hospitals claim to uphold responsible data practices, they have shared patients' identifiable medical data before, sometimes without consent. While Google says that it is "committed to treating that data responsibly and protecting your privacy with strict protocols and innovative privacy technologies," the UK's Information Commissioner's Office ruled that "the Royal Free NHS Foundation Trust failed to comply with the Data Protection Act when it provided patient details to Google DeepMind" (ICO, 2017; Google, n.d.c). Though unsuccessful, a patient sued the University of Chicago Medical Center and Google for sharing patients' medical data that he argued to be not properly anonymized (U.S. District Court, 2020). Ascension shared patients' complete health history, including names and dates of birth, with Google and failed to notify patients or physicians (Copeland, 2019).

Reducing Bias

To reduce bias, inequity, and inequality in AI decision-making, some physicians, patients, and patient advocacies contend that medical AI should use data sets that are more representative and inclusive of the target population. Because "flawed or incomplete data sets that are not inclusive can automate inequality," physician Verghese urges that "if AI is going to make clinicians better at caring for humans in distress, the data sets being used must be representative of society and not biased by sex, race, ethnicity, socioeconomic status, age, ability, and geography" (Thadaney & Verghese, 2018). In addition to one patient warning how "forcing learning when the data isn't there, isn't the right thing to do," another patient says "to make sure you are getting the best training sets from . . . all races" (Dousa, 2020). Since "AI will only produce results as good as the data available to it," the American Hospital Association suggest that "it is critical that organizations have an effective data governance strategy," which includes data quality requirements, such as accuracy, completeness, and relevance (AHA, 2019; Fürber, 2015). Alliance for AI say that "if the representation of the training data used contains features that are not related to the phenomenon being modeled they may also lead to bias" (Allgood et al., 2019). Because "small perturbations in the quality of data input can lead to big mistakes in output," and "data provenance is important to ensure outputs are reliable and applicable even when the precise mechanism by which the data are produced is unknown," physicians Kaushal and Khullar conclude that "careful attention must be paid to both data quality and origins" (Wang et al., 2020).

While insurers and med-tech vendors claim that they take preventative measures to reduce bias, these efforts have not been enough to stop their AI from making biased decisions. Though Optum asserts that "we follow a rigorous process to create algorithms and health care metrics that are unbiased for their intended purpose," they used algorithms to find patients who need more resources, but they determined, falsely, that black patients were less sick (Obermeyer et al., 2019; Optum, n.d.). Microsoft's, IBM's, and Amazon's facial recognition AI were found to be 20.8, 34.4, and 31.4 percent less accurate for dark-skinned women than light-skinned men, respectively (Buolamwini & Gebru, 2018; Raji & Buolamwini, 2019). While Microsoft and IBM were responsive and reduced those inaccuracies to 1.5 and 16.7 percent, respectively, Amazon claimed that the "research paper and article are misleading and draw false conclusions" (Raji & Buolamwini, 2019; Wood, 2019). A year prior to that, Amazon removed an AI recruiting tool that favored male candidates and insisted that the AI "was never used by Amazon recruiters to evaluate candidates," despite men outnumbering women in Amazon's workforce (Dastin, 2018).

Conclusion

Since medical AI has many benefits, physicians should partner with AI such that human judgment and personalized care is balanced with the vast knowledge of AI. However, the concerns of physicians, patients, and patient advocacies should be listened to, which is critical in the early stages when there is pressure to get it right and minimize harmful consequences. Physicians' and patients' distrust of med-tech vendors and the inevitable, increased data collection from AI integration highlight the need to reevaluate patient privacy regulations. Like any technology, AI is not always an appropriate solution, requiring careful analysis to isolate use cases where there is overwhelming proof of its efficacy as well as the least risk of causing harm and deepening social divisions and inequities.

References

Accenture. (2018). *2018 Consumer Survey on Digital Health*. https://www.accenture.com/t20180306T103559Z__w__/us-en/_acnmedia/PDF-71/accenture-health-2018-consumer-survey-digital-health.pdf

Accenture. (2020). *Accenture 2020 Digital Health Consumer Survey*. https://www.accenture.com/_acnmedia/PDF-118/Accenture-2020-digital-health-consumer-survey.pdf

AHA. (2019). *AI and Care Delivery*. https://www.aha.org/center/emerging-issues/market-insights/ai/ai-care-delivery

Allgood, B., Rodriguez, O., Bedorf, J. et al. (2019, September). *Artificial Intelligence in Healthcare: A Technical Introduction*. Alliance for Artificial Intelligence in Healthcare. https://www.theaaih.org/publications/a-primer-on-ai-in-healthcare

Angehrn, Z., Haldna, L., Zandvliet, A. S. et al. (2020). Artificial Intelligence and Machine Learning Applied at the Point of Care. *Frontiers in Pharmacology*, *11*, Article 759. doi:10.3389/fphar.2020.00759

Arnold, M. H. (2021). Teasing Out Artificial Intelligence in Medicine: An Ethical Critique of Artificial Intelligence and Machine Learning in Medicine. *Journal of Bioethical Inquiry*. doi:10.1007/s11673-020-10080-1

Arnold, M. H., Komesaroff, P., & Kerridge, I. (2020). Understanding the Ethical Implications of the Rituals of Medicine. *Internal Medicine Journal*, *50*(9), 1123-1131. doi:10.1111/imj.14990

Bender, E. (2019, December 14). *Unpacking the Black Box in Artificial Intelligence for Medicine*. Undark. https://undark.org/2019/12/04/black-box-artificial-intelligence/

Buolamwini, J., & Gebru, T. (2018). Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification. *Proceedings of Machine Learning Research*, *81*, 77-91. http://proceedings.mlr.press/v81/buolamwini18a.html

Copeland, R. (2019, November 11). Google’s 'Project Nightingale' Gathers Personal Health Data on Millions of Americans. *The Wall Street Journal*. https://www.wsj.com/articles/google-s-secret-project-nightingale-gathers-personal-health-data-on-millions-of-americans-11573496790

Dastin, J. (2018, October 10). *Amazon Scraps Secret AI Recruiting Tool That Showed Bias Against Women*. Reuters. https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G

DeSilva, J., Prensky-Pomeranz, R., & Zweig, M. (2020). *Digital Health Consumer Adoption Report 2020*. Rock Health. https://rockhealth.com/reports/digital-health-consumer-adoption-report-2020/

Dousa, R. (2020). Toward the Clinic: Understanding Patient Perspectives on AI and Data-Sharing for AI-Driven Oncology Drug Development. In Cassidy, J. & Taylor, B. (Eds.), *Artificial Intelligence in Oncology Drug Discovery and Development*, Chapter 8. IntechOpen. doi:10.5772/intechopen.92787

Elwyn, G., Barr, P. J., & Piper, S. (2018). Digital Clinical Encounters. *The BMJ*, *361*, Article k2061. doi:10.1136/bmj.k2061

Farr, C. (2019, March 13). *Cardiologist Eric Topol Explains How Artificial Intelligence Can Make Doctors More Humane*. CNBC. https://www.cnbc.com/2019/03/13/eric-topol-interview-with-christina-farr-on-ai-and-doctors.html

Fürber, C. (2015). *Data Quality Management With Semantic Technologies*. Springer Gabler. doi:10.1007/978-3-658-12225-6

Gao, S., He, L., Chen, Y. et al. (2020). Public Perception of Artificial Intelligence in Medical Care: Content Analysis of Social Media. *Journal of Medical Internet Research*, *22*(7), Article e16649. doi:10.2196/16649

Google. (n.d.a). *Artificial Intelligence at Google: Our Principles*. https://ai.google/principles

Google. (n.d.b). *AI Explanations Whitepaper*. https://storage.googleapis.com/cloud-ai-whitepapers/AI%20Explainability%20Whitepaper.pdf

Google. (n.d.c). *Privacy*. https://health.google/privacy/

IBM. (n.d.). *Explainable AI*. https://www.ibm.com/watson/explainable-ai

ICO. (2017, July 3). *Royal Free - Google Deepmind Trial Failed to Comply With Data Protection Law*. https://ico.org.uk/about-the-ico/news-and-events/news-and-blogs/2017/07/royal-free-google-deepmind-trial-failed-to-comply-with-data-protection-law

Laï, M.-C., Brian, M., & Mamzer, M.-F. (2020). Perceptions of Artificial Intelligence in Healthcare: Findings From a Qualitative Survey Study Among Actors in France. *Journal of Translational Medicine*, *18*, Article 14. doi:10.1186/s12967-019-02204-y

Landi, H. (2019, April 25). *Nearly Half of U.S. Doctors Say They Are Anxious About Using AI-Powered Software: Survey*. Fierce Healthcare. https://www.fiercehealthcare.com/practices/nearly-half-u-s-doctors-say-they-are-anxious-about-using-ai-powered-software-survey

London, A. J. (2019). Artificial Intelligence and Black-Box Medical Decisions: Accuracy Versus Explainability. *The Hastings Center Report*, *49*(1), 15-21. doi:10.1002/hast.973

Longoni, C., Bonezzi, A., & Morewedge, C. K. (2019). Resistance to Medical Artificial Intelligence. *Journal of Consumer Research*, *46*(4), 629-650. doi:10.1093/jcr/ucz013

Matheny, M., Israni, S. T., Ahmed, M. et al. (2019). *Artificial Intelligence in Health Care: The Hope, the Hype, the Promise, the Peril*. National Academy of Medicine. https://nam.edu/artificial-intelligence-special-publication/

Microsoft. (2021). *InterpretML - Alpha Release* (Version 0.2.4). https://github.com/interpretml/interpret

Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). Dissecting Racial Bias in an Algorithm Used to Manage the Health of Populations. *Science*, *366*(6464), 447-453. doi:10.1126/science.aax2342

Optum. (n.d.). *Algorithms and Population Health Management*. https://www.optum.com/business/health-insights/algorithms-human-touch.html

Pearl, R. (2019, December 16). *Why Big Tech Companies Won't Solve Healthcare's Biggest Challenges*. Forbes. https://www.forbes.com/sites/robertpearl/2019/12/16/big-tech/?sh=266608c16d28

Promberger, M., & Baron, J. (2006). Do Patients Trust Computers? *Journal of Behavioral Decision Making*, *19*(5), 455-468. doi:10.1002/bdm.542

PwC. (2017, June). *What Doctor? Why AI and Robotics Will Define New Health*. https://www.pwc.com/gx/en/industries/healthcare/publications/ai-robotics-new-health/survey-results.html

Raji, I. D., & Buolamwini, J. (2019, January 24). *Actionable Auditing: Investigating the Impact of Publicly Naming Biased Performance Results of Commercial AI Products*. MIT Media Lab. https://www.media.mit.edu/publications/actionable-auditing-investigating-the-impact-of-publicly-naming-biased-performance-results-of-commercial-ai-products/

Sgaier, S. K., Huang, V., & Charles, G. (2020). *The Case for Causal AI*. Stanford Social Innovation Review. https://ssir.org/articles/entry/the_case_for_causal_ai

Szabo, L. (2019, December 24). *Artificial Intelligence Is Rushing Into Patient Care—and Could Raise Risks*. Scientific American. https://www.scientificamerican.com/article/artificial-intelligence-is-rushing-into-patient-care-and-could-raise-risks/

Teredesai, A., Eckert, C., Ahmad, M. A., & Kumar, V. [Association for Computing Machinery]. (2018, October 12). *Explainable Machine Learning Models for Healthcare AI* [Video]. YouTube. https://www.youtube.com/watch?v=4pgLsDzrlB8

Thadaney, S., & Verghese, A. (2018). Humanizing Artificial Intelligence. *JAMA*, *321*(1), 29-30. doi:10.1001/jama.2018.19398

Thingalaya, N. K. (2019, October 14). *Why This Doctor Thinks We Need More Transparency for Algorithms in Health Care*. Op-Med. https://opmed.doximity.com/articles/why-this-doctor-thinks-we-need-more-transparency-for-algorithms-in-health-care

Tonekaboni, S., Joshi, S., McCradden, M. D., & Goldenberg, A. (2019). What Clinicians Want: Contextualizing Explainable Machine Learning for Clinical End Use. *Proceedings of Machine Learning Research*, *106*, 359-380. http://proceedings.mlr.press/v106/tonekaboni19a.html

Topol, E. J., & Verghese, A. (2020, January 2). *Bringing Humanity Back Is the 'Greatest Gift AI Can Give'*. Medscape. https://www.medscape.com/viewarticle/922600

Topol, E. J., Verghese, A., & Ofri, D. (2019, August 22). *AI May Save Some Time, but It Can't Listen to a Patient*. Medscape. https://www.medscape.com/viewarticle/916392

Tran, V.-T., Riveros, C., & Ravaud, P. (2019). Patients' Views of Wearable Devices and AI in Healthcare: Findings From the Compare E-Cohort. *npj Digital Medicine*, *2*, Article 53. doi:10.1038/s41746-019-0132-y

U.S. District Court for the Northern District of Illinois: Eastern Division. (2020). Dinerstein v. Google, LLC et al., 85 Il. 1079 (2020). https://law.justia.com/cases/federal/district-courts/illinois/ilndce/1:2019cv04311/366172/85/

Verghese, A. (2011, February 26). Treat the Patient, Not the CT Scan. *The New York Times*. https://www.nytimes.com/2011/02/27/opinion/27verghese.html

Verghese, A., Shah, N. H., & Harrington, R. A. (2018). What This Computer Needs Is a Physician. *JAMA*, *319*(1), 19-20. doi:10.1001/jama.2017.19198

Wang, F., Kaushal, R., & Khullar, D. (2019). Should Health Care Demand Interpretable Artificial Intelligence or Accept "Black Box" Medicine? *Annals of Internal Medicine*, *172*(1), 59-60. doi:10.7326/M19-2548

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

Watson, D. S., Krutzinna, J., Bruce, I. N. et al. (2019). Clinical Applications of Machine Learning Algorithms: Beyond the Black Box. *BMJ*, *364*(886). doi:10.1136/bmj.l886

Wood, M. (2019, January 26). *Thoughts on Recent Research Paper and Associated Article on Amazon Rekognition*. Amazon. https://aws.amazon.com/blogs/machine-learning/thoughts-on-recent-research-paper-and-associated-article-on-amazon-rekognition/

Wronski, L. (2020, June 26). *Fortune|SurveyMonkey poll: big tech consolidation*. SurveyMonkey. https://www.surveymonkey.com/curiosity/fortune-poll-big-tech/

York, T., Jenney, H., & Jones, G. (2020). Clinician and Computer: A Study on Patient Perceptions of Artificial Intelligence in Skeletal Radiography. *BMJ Health & Care Informatics*, *27*, Article e100233. doi:10.1136/bmjhci-2020-100233