Articles   Faculty Scholarship

2019

Artificial Intelligence in the Medical System: Four Roles for Potential Transformation

Will Nicholson Price II

University of Michigan Law School, wnp@umich.edu

Recommended Citation

Artificial Intelligence in the Medical System: Four Roles for Potential Transformation

W. Nicholson Price II[*]

ABSTRACT

Artificial intelligence (AI) looks to transform the practice of medicine. As academics and policymakers alike turn to legal questions, a threshold issue involves what role AI will play in the larger medical system. This Article argues that AI can play at least four distinct roles in the medical system, each potentially transformative: pushing the frontiers of medical knowledge to increase the limits of medical performance, democratizing medical expertise by making specialist skills more available to non-specialists, automating drudgery within the medical system, and allocating scarce medical resources. Each role raises its own challenges, and an understanding of the four roles is necessary to identify and address major hurdles to the responsible development and deployment of medical AI.

Artificial intelligence (AI) is already beginning to change the practice of medicine—but how will that change manifest? Medical AI typically uses newly available troves of health and health-related data to find underlying patterns that can be used to make recommendations or decisions. Its increasing use raises plenty of legal and ethical issues such as informed consent and privacy, considered in other articles in this Symposium.[1] In addition to questions about how medical AI functions on its own, or how it is developed, questions arise about how AI will interact with the health system more generally. Some have worried that AI will "replace" physicians,[2] but more pressing concerns involve how AI and clinicians will interact, when clinicians should override AI recommendations even if they disagree, and how AI will alter the daily routine of clinicians and patients.[3] Tackling these issues requires understanding what AI actually does and can do.

One way to think about medical AI is how it does what it does: neural networks versus support-vector machines, explainable algorithms versus black-box methods, and similar categories.[4] Another is to categorize the specific tasks AI can accomplish: diagnostics, prognostics, image analysis, treatment recommendations, or similar specific tasks.[5] But these analyses elide something about the broader picture of AI in medicine.

This piece considers a different typology: what role will AI play in the medical system?[6] It suggests that there are at least four roles that AI can play, each of which

---

1. *See, e.g.*, Nathan Cortez, *Digital Health and Regulatory Experimentation at the FDA*, 18 YALE J. HEALTH POL'Y L. & ETHICS 4 (2019), 21 YALE J.L. & TECH. 4 (2019); A. Michael Froomkin, *Big Data: Destroyer of Informed Consent*, 18 YALE J. HEALTH POL'Y L. & ETHICS 27 (2019), 21 YALE J.L. & TECH. 27 (2019); Nicolas Terry, *Of Regulating Healthcare AI and Robots*, 18 YALE J. HEALTH POL'Y L. & ETHICS 133 (2019), 21 YALE J.L. & TECH. 133 (2019).

2. *See, e.g.,* Liat Clark, *Vinod Khosla: Machines Will Replace 80 Percent of Doctors,* WIRED (Sep. 4, 2012), https://www.wired.co.uk/article/doctors-replaced-with-machines; Hannah Devlin, *London Hospitals to Replace Doctors and Nurses with AI for Some Tasks*, GUARDIAN (May 21, 2018), https://www.theguardian.com/society/2018/may/21/london-hospitals-to-replace-doctors-and-nurses-with-ai-for-some-tasks; *but see* Will Knight, *Google X-ray Project Shows AI Won't Replace Doctors Any Time Soon*, MIT TECH. REV. (March 27, 2018), https://www.technologyreview.com/s/610552/google-x-ray-project-shows-ai-wont-replace-doctors-any-time-soon; Marschall S. Runge, *Artificial Intelligence Will Put a Premium on Physicians' Knowledge and Decision-making Skills*,

3. *See, e.g.*, Federico Cabitza, Raffaele Rasoini & Gian Franco Gensini, *Unintended Consequences of Machine Learning in Medicine*, 318 JAMA 517, 517–18 (2017).

4. *See*, *e.g.*, Fei Jiang et al., *Artificial Intelligence in Healthcare: Past, Present and Future*, 2 STROKE AND VASCULAR NEUROLOGY 230 (2017) (providing an overview of medical AI mechanisms); Will Knight, *The Dark Secret at the Heart of AI*, MIT TECH. REV. (Apr. 11, 2017), https://www.technologyreview.com/s/604087/the-dark-secret-at-the-heart-of-ai/ (addressing explainable versus black-box algorithms)

5. *See, e.g.,* Ziad Obermeyer & Ezekiel J. Emanuel, *Predicting the Future—Big Data, Machine Learning, and Clinical Medicine*, 375 NEW ENG. J. MED. 1216, 1217 (2016).

6. AI can also play roles in the health system outside the medical system; for instance, AI can help find epidemiological patterns or identify adverse reactions to drugs that are used in the market.

has the potential to transform medical practice in a different way. First, AI can push the frontiers of medical knowledge and practice, allowing forms of care that were previously unavailable. Second, AI can democratize medical expertise, allowing broader groups of clinicians the ability to perform tasks previously limited to subsets of specialists, but well within existing human capacity. Third, AI can automate drudgery, performing quotidian tasks such as paperwork and freeing clinician time for other tasks. Fourth, AI can allocate scarce resources, directing them where they can do the most good.

It is easy to focus on one or two of AI's potential roles: the headline promise of AI is to push medical frontiers, and the most straightforward forms of AI to market aim to democratize expertise and replicate existing capabilities. Indeed, when discussions focus on the role of the Food and Drug Administration (FDA) as a quality gatekeeper for medical AI promotes consideration of democratizing expertise and, to a lesser extent, pushing frontiers.[7] But each role raises different challenges for development and deployment, and as academics, developers, and policymakers consider how AI can best be integrated into a changing medical system, all of these four roles should be part of that analysis.

I. FOUR ROLES FOR MEDICAL AI

Each of the four roles AI can serve in the medical system has the potential to transform medical practice in the relatively near term; this Section describes these four roles and gives examples of each.

A. Pushing frontiers

The most hyped and most exciting promise of medical AI is the possibility of pushing the frontiers of medical knowledge and medical care, letting experts accomplish otherwise-impossible medical feats. Such care can be more precise, and better tailored to the individual patient, than what existing explicit medical knowledge allows, especially where the right answer depends on complex networks of underlying factors. For instance, interactions between thousands of different genes can help predict a tumor's response to different chemotherapeutics. These interaction patterns may be too complex for explicit models, but AI can use them to recommend the choice of drugs[8] or to predict patient responsiveness to

---

These uses are outside the scope of this Article, which focuses on the medical system.

7. *See, e.g.*, W. Nicholson Price II, *Regulating Black-Box Medicine*, 116 MICH. L. REV. 421 (2017) (considering FDA regulation of medical AI).

8. *See, e.g.*, Michael Q. Ding et al., *Precision Oncology beyond Targeted Therapy: Combining Omics Data with Machine Learning Matches the Majority of Cancer Cells to Effective Therapeutics*, 16 MOLECULAR CANCER RES. 269, 269 (2017) (using genomic data to predict drug

treatment.[9] AI can similarly use a broad set of factors to predict the likelihood of a patient dying in the next year with high accuracy which can enable different choices about treatment or life planning.[10]

AI can also enable continuously tailored treatment of the sort otherwise impossible today. For patients with diabetes, "artificial pancreases" are under development—and the first are newly available—which combine a controller with a subcutaneous insulin pump and a continuous glucose monitor.[11] Using artificial intelligence and machine learning techniques, the controller can "learn" the relationship between insulin dosages, other relevant information such as time of day, and glucose levels.[12] It can then adjust the small dosages of insulin received throughout the day to improve patient blood sugar levels more precisely than either explicit algorithms or periodic patient adjustments.[13]Ideally, the system continuously learns, starting from a relatively standardized model and adjusting itself to the particular patient the system is aiding.

In this flagship model, then, AI augments the capacity of even the best human doctors and most capable health systems, allowing them to do what they could not before.

B. Replicating and democratizing expertise

AI can also democratize medical expertise—that is, it can make expertise available more widely than it currently is. AI can augment the capacity of generalists, so they can accomplish tasks currently limited to specialists, and it can also provide at least some care in situations with very limited current availability of any care at all.

Medical practice is limited by the availability (and cost) of specialist providers, and AI can augment the abilities of generalists to decrease that limitation, enabling more patients to access specialist care. Many patients cannot see ophthalmologists when they should—for instance, for yearly diabetic

---

effectiveness in cancer cell lines).

9. *See, e.g.*, Nathaniel Braman et al*., Predicting Neo-adjuvant Chemotherapy Response From Pre-treatment Breast MRI Using Machine Learning and HER2 Status,* 36 J. CLIN. ONCOL. 582 (supp. 15) (2018) (using imaging and status of the HER2 gene to predict response to neo-adjuvant chemotherapy).

10. Alvin Rajkomar et al., *Scalable and Accurate Deep Learning with Electronic Health Records*, 1 NPJ DIGITAL MED., no. 18 (2018).

11. *See, e.g.*, V. Buch, G. Varughese, & M. Maruthappu, *Artificial Intelligence in Diabetes Care*, 35 DIABETIC MED. 495 (2018).

12. *Id.*

13. W. Nicholson Price II, *We Have to Be Smart About Artificial Intelligence in Medicine*, SLATE (Aug. 15, 2018, 9:00 AM), https://slate.com/technology/2018/08/a-i-in-diabetes-care-can-we-trust-it-if-we-dont-know-how-it-works.html.

retinopathy screens for patients with diabetes—in part because there are too few ophthalmologists, or because ophthalmologist visits are too expensive or otherwise inaccessible for some patients.[14] This is especially true in the developing world.[15] AI can help here also; for instance, the FDA approved the IDx-DR system in May 2018 for autonomous diagnosis of diabetic retinopathy.[16] The IDx-DR consists of an automated camera that takes pictures of the retina; an operator with minimal training uses the camera, which transmits images to one algorithm for quality and then another algorithm, which determines if the patient has more-than-mild diabetic retinopathy, in which case the patient should seek further care.[17] IDx-DR can allow a generalist, such as a primary-care physician or nurse-practitioners, to provide screening care that would otherwise require an ophthalmologist.[18] Similarly, Google is developing AI-driven ophthalmological programs for deployment in India.[19] Along the same lines, algorithms can already identify malignant skin lesions as well as dermatologists in some situation;[20] patients with positive results would then presumably undertake the additional effort to seek out dermatologists or oncologists for treatment.

AI can also bring care to patients in situations where provider access is extremely limited or nonexistent. In rural areas of India, one primary care physician has responsibility for 30,000 patients.[21] Liberia had six ophthalmologists in 2016 for a population of 4.5 million.[22] Rural areas in the U.S. likewise face persistent physician shortages, though nowhere near as severe.[23] AI can give

14. Karinya Lewis, *Improving Patient Compliance with Diabetic Retinopathy Screening and Treatment*, 28 COMM. EYE HEALTH 68 (2015).

15. *See, e.g.*, Alfred Sommer et al., *Challenges of Ophthalmic Care in the Developing World*, 132 JAMA OPHTHALMOLOGY 640 (2014).

16. U.S. FDA, *FDA Permits Marketing of Artificial Intelligence-Based Device to Detect Certain Diabetes-related Eye Problems* (Apr. 11, 2018), https://www.fda.gov/newsevents/ newsroom/pressannouncements/ucm604357.htm.

17. IDx-DR, IDX, https://www.eyediagnosis.net/idx-dr.

18. *Id.*

19. Tom Simonite, *Google's AI Eye Doctor Gets Ready to Go to Work in India*, WIRED (June 8, 2017, 2:25 PM), https://www.wired.com/2017/06/googles-ai-eye-doctor-gets-ready-go-work-india; Varun Gulshan et al., *Development and Validation of a Deep Learning Algorithm for Detection of Diabetic Retinopathy in Retinal Fundus Photographs*, 316 JAMA 2402 (2016).

20. Holger A. Haenssle et al., *Man Against Machine: Diagnostic Performance of a Deep Learning Convolutional Neural Network for Dermoscopic Melanoma Recognition in Comparison to 58 Dermatologists*, 29 ANN. ONCOLOGY 1836, 1836 (2018).

21. Devarsetty Praveen et al., *SMARTHealth India: Development and Field Evaluation of a Mobile Clinical Decision Support System for Cardiovascular Diseases in Rural India*, 2 JMIR MHEALTH & UHEALTH e54 (2014).

22. Al-Varney Rogers, *Liberian Doctors Threaten Go-Slow over Salary Arrears*, FRONT PAGE AFRICA (Nov. 15, 2016), https://frontpageafricaonline.com/health/liberian-doctors-threaten-go-slow-over-salary-arrears/.

23. Amelia Goodfellow et al., *Predictors of Primary Care Physician Practice Location in*

patients direct access to certain forms of medical care (including diagnosis or basic recommendations, for instance), and can similarly triage those patients who do in fact require access to whatever providers are available. "Ping An Good Doctor," a Chinese platform, plans to combine a small waiting room about the size of a phone booth with an interactive AI-based system that solicits patient information, builds a medical history, and then communicates the relevant information to a remote cloud-based physician.[24] The patient would then be able to purchase medicine from the adjoining "Smart Medicine Cabinet" or a nearby pharmacy.[25] Such systems allow access to care for patients who might otherwise have little or no such access.

C. Automating drudgery

AI has the potential to transform daily medical practice in a more straightforward way by automating medical drudgery. This automation could free substantial amounts of clinician time, profoundly changing the experience of medical practice and the ability of clinicians to interact with patients more deeply.[26]

Certain types of automation can streamline the attention that clinicians pay to some forms of data or some patients, along the lines of the triage functions described above. If AI can triage radiologic images and identify those which can safely be ignored, those with a high probability of problems, and those about which the AI is uncertain, radiologists can focus on the latter two and save themselves a substantial amount of work.[27]

AI can also transform the much-maligned experience of clinician. Physicians spend almost half of their time working with electronic health records and other

---

Underserved Urban and Rural Areas in the United States: A Systematic Literature Review, 91 ACAD. MED. 1313 (2016).

24. *The First "Unmanned Clinic" Reveals Itself in Wuzhen; Ping An Good Doctor Uses AI Technology to Empower Healthy City*, PRNEWSWIRE (Nov. 7, 2018), https://en.prnasia.com/releases/apac/the-first-unmanned-clinic-reveals-itself-in-wuzhen-ping-an-good-doctor-uses-ai-technology-to-empower-a-healthy-city-228380.shtml.

25. *Id.*

26. *See* Greg Irving et al., *International Variations in Primary Care Physician Consultation Time: A Systematic Review of 67 Countries*, 7 BMJ OPEN e017902, 5–10 (2017) (documenting short primary care consultation times across many countries). Of course, the idea that automating drudgery would increase patient interaction time assumes that given additional time, clinicians will choose to spend more time interacting with the same patients, rather than just increasing the volume of patients they see. This assumption may not be warranted, though clinicians profess to want to spend more time with patients. *See, e.g.*, Chris Hayhurst, *What Doctors Want: More Time with Patients*, ATHENAINSIGHT (Oct. 18, 2017), https://www.athenahealth.com/insight/what-doctors-want-more-time-patients.

27. Thomas H. Davenport & Keith J. Dreyer, *AI Will Change Radiology, but It Won't Replace Radiologists*, HARV. BUS. REV. (Mar. 27, 2018), https://hbr.org/2018/03/ai-will-change-radiology-but-it-wont-replace-radiologists.

desk work (such as data entry and data search).[28] Approximately one third of physician time in the examination room itself is devoted to such tasks.[29] This large amount of non-care work is cited as the leading cause of physician burnout and is a source of perennial complaints.[30] AI could automate at least some fraction of paperwork. At the front end of the patient encounter, AI can help with routine search functions, identifying relevant data from patient records or from the medical literature. During and after the patient encounter, AI-driven natural language processing and speech recognition could act as a scribe, taking notes rather than requiring the clinician to input notes manually; ideally, the AI could even identify the most relevant points and structure the data for entry into electronic health records.[31] Such interventions could save clinician time and increase the time available for direct interactions with patients as human beings—something AI largely can't do. In addition to saving time, AI note-taking could at least potentially improve the accuracy of notes by reducing errors endemic in human-run data entry.[32]

D. Allocating resources

Finally, AI can guide the allocation of scarce resources available within the medical system. The earlier roles for medical AI focused on either extending the capabilities of clinicians, creating resources where none were available, or freeing up additional clinician time. AI can also help ensure that whatever scarce resources exist—and even with the above interventions, there will still likely be some forms of scarcity—are allocated to the patients that need them most, or that can benefit from them most. AI could thus allocate drugs during a shortage[33] or inpatient hospitals beds when more patients need beds than are available.[34] On a shorter

---

28. Christine Sinsky et al., *Allocation of Physician Time in Ambulatory Practice: A Time and Motion Study in 4 Specialties*, 165 ANN. INTERNAL MED. 753, 753 (2016).

29. *Id.*

30. *See, e.g.*, Sandhya K. Rao, et al., *The Impact of Administrative Burden on Academic Physicians: Results of a Hospital-wide Physician Survey*, 92 ACAD. MED. 237, 237–43 (2017) (finding "[p]hysicians who reported higher percentages of time spent on administrative duties had . . . higher levels of burnout")

31. Enrico Coiera et al., *The Digital Scribe*, 1 NPJ DIGITAL MED. art. 58 (2018); Steven Y. Lin, Tait D. Shanafelt & Steven M. Asch, *Reimagining Clinical Documentation with Artificial Intelligence*, 93 MAYO CLIN. PROC. 593 (2018).

32. Reduction in error depends on AI itself not introducing errors, or at least introducing fewer errors than the human activity it replaces.

33. *See, e.g.*, Matthew DeCamp et al., *Chemotherapy Drug Shortages in Pediatric Oncology: A Consensus Statement*, 133 PEDIATRICS peds-2013 (2014) (describing the problem of chemotherapy drug allocation during shortages and noting the ethical complexity of the issue).

34. I. Glenn Cohen et al., *The Legal and Ethical Concerns that Arise from Using Complex Predictive Analytics in Health Care*, 33 HEALTH AFF. 1139, 1140 (2014).

timescale, AI could allocate the scarce resources of the emergency room, triaging patients and stratifying them by risk; several algorithms have also been developed to assist with emergency radiology, especially for head injuries.[35] Implicit within such decisions are a range of ethical issues about desert and allocation; should resources be given to the first patients to present, the patients that need them the most, the patients that could benefit most, or the patients that can pay the most?[36] Medical ethicists, hospital administrators, and insurance payers might not agree on the appropriate quantity to optimize, and indeed whoever develops the AI might not make that optimization readily knowable. In any case, it seems likely that AI will be involved in hospital or health-system resource allocation (and, later, billing), whether AI's presence is visible to the patient population or largely remains behind the scenes.

II. ROLE-SPECIFIC CHALLENGES

The variety of roles that AI can play in the medical system leads to a similar variety of challenges in ensuring high-quality AI applications. This Article does not attempt a comprehensive listing or analysis of those challenges, but instead offers a cautionary note: It is seductively easy to focus on FDA as locus of quality control, because FDA plays such a role for many medical technologies, and medical AI is indeed a medical technology. And FDA may be well-suited to addressing some types of quality challenges, as I have argued before,[37] like evaluating whether the AI accurately predicts an outcome or suggests an effective outcome in clinical trials. But understanding the systemic roles AI can play leads to identifying other challenges that spring from those roles.

To take one example, democratizing medical expertise raises the challenge of overcoming contextual bias, so that AI actually provides high-quality medical information to patients in different medical contexts. Medical AI, at least in the U.S., is typically developed in a rather specific subset of environments, namely those with relatively high resources such as academic medical centers. Machine learning techniques learn from the data they have available, including the data about patients, about medical practices and procedures, and about outcomes (when such data are available). But the patients seen in these high-resource medical environments generating these data differ systematically from those seen in low-resource environments. For instance, in New York City, academic medical centers see systematically fewer African American patients than other hospitals.[38] More

---

35. Yosef Berlyand et al., *How Artificial Intelligence Could Transform Emergency Department Operations*, 36 AM. J. EMERGENCY MED. 1515, 1515 (2018).

36. *Id.*

37. Price, *supra* note 7.

38. Roosa Sofia Tikkanen et al., *Hospital Payer and Racial/Ethnic Mix at Private Academic*

generally, medical big data has substantial biases; males, individuals of European descent, and wealthier patients are overrepresented in the data and especially in the data used to train AI.[39] These non-representative data may substantially skew the diagnostic and therapeutic patterns that AI sees, resulting in lower-accuracy diagnoses and therapy recommendations for patients from different populations.

In addition to differences between populations, patterns of care also differ substantially between environments. These differences may be based on the variation inherent in the health-care system or may result from reasonable differences in what the benefits and costs are of various treatments considering the resources available. For instance, otherwise riskier interventions may be acceptable when resources are available for mitigating the consequences of those risks. In either case, AI trained on those data will tend to incorporate the care patterns of the training environment. Thus, whether AI makes diagnoses or recommendations, they may be heavily dependent on the environments in which they were developed. As a result, AI-recommended care suggested to patients outside those environments may be less relevant or effective in their care setting. FDA can potentially address such challenges by requiring evidence of cross-context performance, but recognizing the need for such evidence requires acknowledging the goal of democratizing expertise in the first place.

But that is just one role-specific problem, and one which fits relatively readily into FDA's existing paradigm of demonstrating quality. Other roles for medical AI bring their own problems. Pushing medical frontiers raises challenges of identifying ground truth: When AI performs at a level beyond what humans can match, how do we know when it's working? Robust clinical trials can certainly help, but clinical trials are not always feasible, especially for algorithms that ideally learn over time as they are exposed to new data. Observing outcomes over time in the context of a learning health system can help establish ground truth more readily over time, but this demands a different validation model than FDA traditionally uses.

Automating drudgery raises issues entirely outside FDA's purview. Software that simply interacts with patient records (rather than making an individual recommendation or diagnosis) is categorically excluded from the definition of medical devices by the 21st Century Cures Act.[40] But low-quality voice-recognition or data-input software can taint not only individual patient records, but also the

---

Medical Centers in Boston and New York City, 47 INT'L J. HEALTH SERVS. 460, 464 (2017).

39. Dave Gershgorn, *If AI is Going to be the World's Doctor, It Needs Better Textbooks,* QUARTZ (Sept. 6, 2018), https://qz.com/1367177/if-ai-is-going-to-be-the-worlds-doctor-it-needs-better-textbooks; Alice B. Popejoy & Stephanie M. Fullerton, *Genomics Is Failing on Diversity*, 538 NATURE 161 (2016).

40. 21st Century Cures Act § 3060, 130 Stat. at 1130–33.

data on which other validation efforts may rest.

Finally, resource allocation AI raises challenges even further afield. Such AI does not fall within the Act's definition of medical devices, largely removing FDA's ability to regulate it.[41] And the quality questions involved require far more than validating that the AI does what it says it does, because resource allocation implicates fundamental ethical questions of justice and equity in access to scarce medical resources. If a cost-focused AI recommends delaying a clinician's visit to see an insured patient, that could be because other patients are in more urgent need of care—or because the AI can predict that the first patient will die rapidly without a clinician visit, thus saving the hospital a bundle in uncompensated care. That outcome is tremendously problematic but has nothing to do with whether the AI is working, and instead turns on the objectives the AI is given.

CONCLUSION

AI will transform the practice of medicine, but that transformation will not be uniform. In addition to the different ways AI can function in and of itself, policymakers, academics, and developers should take into account the role any particular AI instantiation will play in the medical system as a whole. Each role— pushing frontiers, democratizing expertise, automating drudgery, and allocating resources—raises different quality challenges, and requires its own care to ensure that medical AI transforms medicine for the better.

---

41. *Id.*