

The Crucial Role of Academic Medical Centers in Preparing for the Future of Health AI

Michael J. Pencina, Mary E. Klotman

A global technological revolution driven by artificial intelligence (AI) is unfolding before our eyes, a revolution further accelerated by the recent public debut of AI platforms incorporating large language models. Although the ultimate effects remain to be seen, there's little doubt that AI will bring substantial changes to a myriad of different sciences and industries, including health care. Indeed, AI is already being applied to a spectrum of health tools spanning diagnosis, treatment, and the management of care delivery. AI applications are also being incorporated into basic, clinical, and translational research and used in patient-provider communication and medical education.

However, as the pace of AI development and uptake accelerates, these technologies—many of them relatively untested—present challenges for the national health care enterprise. Academic medical centers (AMCs) are struggling to keep up with the breakneck pace of change and navigate an ever-increasing resource mismatch between academia and industry. At the same time, AI tools have the potential to enhance the success of our core missions while mitigating this resource mismatch. AMCs have an unprecedented opportunity to forge thoughtful partnerships and assume a leadership role in the responsible implementation of these new, powerful tools. However, this will require AMCs to reimagine and adapt their approaches to their core missions of care delivery, research, education, and community partnerships to achieve new goals while ensuring the safety and well-being of patients. It is thus imperative for AMCs to adopt a strategic approach that identifies priority areas for growth and development and enables expansion by building on existing strengths. In this article, we share our experience embarking on this journey at the Duke University School of Medicine and Duke Health.

Ensuring Trustworthy Health AI in Clinical Care and Research

Academia initially played a leading role in the development of AI, starting with its emergence as a field of study in

the 1950s [1]. More recently, however, the center of gravity in AI development has shifted and now clearly resides with industry. A recent report underscores this gap, with industry releasing 51 “notable” machine learning models over the past 20 years compared with 15 for academia and 21 for academic-industry collaborations [2].

However, while industry can invest vast resources and expertise in new AI technologies, academia has a critical role to play in ensuring that what is being developed is trustworthy and has a greater societal benefit as its ultimate goal. Many AI tools have substantial potential, but the complexity of their inner workings may mask flaws and biases that only manifest with “real-world” use [3–5]. In health AI, enthusiasm for new technologies and “fear of missing out” have often converged to create a “Wild West” environment. In such an atmosphere, the impetus to rapidly adopt new technologies may replace careful and cautious evaluation. If AI has the potential to enhance clinical decision-making, then as with any other new diagnostic or therapeutic technology, an evidence-based approach to validating and monitoring its application to clinical care is warranted. For this reason, there is a pressing need for multidisciplinary expertise that can provide a thorough, impartial assessment of AI technologies throughout their lifecycles.

At Duke, we have invested in programs and initiatives aimed at ensuring that AI tools developed or adopted for use in our health care environments are subjected to stringent, methodical scrutiny not just of the tool's technical attributes, but also for issues of ethics, equity, and fairness that may arise from its use. The mission of Duke AI Health, one of our first such programs, evolved from fostering the development of AI tools to focusing on enabling ethical and equitable data science across the entire health care and research enterprise. A key initiative arising from this focus was Duke's Algorithm-Based Clinical Decision Support (ABCDS) Oversight. Focused on people, process, and technology, it grew as a partnership between

our clinical and research missions in order to ensure that any algorithm deployed at Duke for patient care or clinical research was carefully evaluated for safety, effectiveness, and equity before its release, and its performance monitored continuously throughout its lifecycle [6].

More recently, Duke has broadened its efforts as a co-founder of the national Coalition for Health AI (CHAI) initiative, which engages in large-scale efforts to ensure trustworthy health AI for patient benefit, including a recently published Assurance Standards Guide and reporting checklist for developing and deploying AI in health care [7]. Duke and other AMCs have also engaged with industry partners such as Microsoft through the Trustworthy & Responsible AI Network (TRAIN), which seeks to research and develop tools that facilitate the application and monitoring of AI solutions, with patient safety and benefit as the top priority [8].

Education and Workforce Development

Success in efforts to build trustworthy health AI depends on a skilled and knowledgeable workforce capable of using AI tools responsibly in appropriate contexts—and just as importantly, understanding their limitations [9]. An AI-literate workforce is essential for implementing AI-based health technologies in ways that ensure benefit and value for patients, families, health care providers, and society. This needs to be accomplished on different levels depending on competencies required for different roles. For example, frontline clinicians who use ambient voice transcription technologies should be trained in their use, with sufficient emphasis on learning about undesirable or unintended consequences. Furthermore, the proliferation of chatbots—computer programs that simulate human conversation—will require clinicians to be ready for new kinds of questions and concerns from patients. On a larger scale, as AI tools are integrated into the daily tasks of medical care, we will need to reimagine our approach to medical education [10]. Providers need to be informed of the tools being used and know how to actively monitor and engage with these tools in an ethical manner.

These needs cannot be fulfilled solely by industry, which tends to prioritize the creation of products for the market. However, they do dovetail closely with the core missions of AMCs and their universities. There is an opportunity for mutually beneficial partnerships with industry in

which AMCs provide critical context and education that informs the development of these tools as well as generating scholarly evidence that informs responsible and safe implementation. Multiple programs are underway at Duke to build capacity for an AI-literate workforce, including the AI Health Fellowship Program, an intensive two-year training program in health data science that embeds learners with backgrounds in quantitative sciences within multidisciplinary teams that include clinical and statistical experts. More broadly, the AI Health Seminar Series provides access to experts in AI, machine learning, and related fields through lectures and workshops, many of which are free and open to the public.

Most recently, Duke AI Health, in conjunction with the Duke School of Nursing, helped serve as an incubator for the Fostering AI/ML Research for Health Equity and Learning Transformation Hub (FAIR HEALTH) program, which provides a central resource for equipping nurses with the key competencies needed for using AI-based tools while ensuring safe and equitable care for patients.

Informing Policy

AMCs can make substantive contributions to trustworthy health AI in the domain of policy. The Duke Margolis Institute for Health Policy has been engaged in multiple partnerships and crosscutting projects aimed at informing regulatory challenges and solutions related to the implementation of health AI and the data underpinning its use. In addition, the Duke Clinical Research Institute has developed a portfolio of projects with a specific focus on leveraging AI applications to support transformation of the clinical research enterprise. This includes a center devoted to re-engineering clinical trials to take advantage of AI capabilities, digital health technologies, novel data sources, and emerging methods for conducting research that can inform practice and improve access to equitable health care.

Patient and Community Engagement

Finally, we recognize that health AI cannot progress unless patients, families, and communities are engaged as equal partners. AI tools must be implemented and reported on transparently to engender widespread trust. A critical element of such trust is ensuring that the patient voice is part of development and deployment of health AI.

The world of AI is filled with enthusiasm, often genu-

ine, about the potential benefits these technologies offer. However, uncritical enthusiasm can lead to biased hype. Trusted, impartial partners are needed to bridge the gap between the AI community and a public that may harbor reasonable skepticism about the usefulness, safety, and fairness of the products it generates. We believe that AMCs can be a crucial bridge between health AI innovators and their intended beneficiaries: patients, clinicians, and communities.

Conclusion

AI applications are already bringing transformative change to health care delivery and clinical research. However, in order for these technologies to realize their potential for benefit, clinicians, health systems, payers, patients, and the public must have confidence that they are being developed and used in ways that are transparent, equitable, and beneficial. AMCs are poised to play a central role in ensuring the trustworthiness of algorithmic technologies in health care. NCMJ

Michael J. Pencina, PhD Chief Data Scientist, Duke Health, Vice Dean for Data Science and Director, Duke AI Health, Duke University School of Medicine, Durham, North Carolina.

Mary E. Klotman, MD Executive Vice President for Health Affairs, Duke University; Dean, Duke University School of Medicine; Chief Academic Officer, Duke Health, Durham, North Carolina.

Acknowledgments

The authors thank Jonathan McCall, MS, for editorial assistance in the preparation of this manuscript.

Dr. Pencina is a co-founder and board member of the Coalition for Health AI (CHAI). He also reports serving as a consultant for RevelAI and has received funding from the Moore Foundation for assessment of health AI maturity. Dr. Klotman has no disclosures related to this work to report.

References

1. Moor J. The Dartmouth College Artificial Intelligence Conference: The Next Fifty Years. *AI Magazine*. Published 2006. Accessed April 18, 2024. <https://aaai.org/ojs/index.php/aimagazine/article/view/1911/1809>
2. Aiken C, Auricchio T, Besiroglu T, et al. Chapter 1: Research and Development. In: *Measuring trends in AI: Artificial intelligence index report*, 2024. Perrault R, Clark J, Eds. Stanford University Human-Centered Artificial Intelligence; 2024. Accessed April 18, 2024. <https://aiindex.stanford.edu/report/>
3. Rudin C. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nat Mach Intell*. 2019;1(5):206-215. doi: 10.1038/s42256-019-0048-x
4. Vela D, Sharp A, Zhang R, Nguyen T, Hoang A, Panykh OS. Temporal quality degradation in AI models. *Sci Rep*. 2022;12(1):11654. doi: 10.1038/s41598-022-15245-z
5. Lyons PG, Hafford MR, Yu SC, et al. Factors associated with variability in the performance of a proprietary sepsis prediction model across 9 networked hospitals in the US. *JAMA Intern Med*. 2023;183(6):611-612. doi: 10.1001/jamainternmed.2022.7182
6. Economou-Zavlanos NJ, Bessias S, Cary MP Jr, et al. Translating ethical and quality principles for the effective, safe and fair development, deployment and use of artificial intelligence technologies in healthcare. *J Am Med Inform Assoc*. 2024;31(3):705-713. doi: 10.1093/jamia/ocad221
7. Assurance Standards Guide and Reporting Checklist. Coalition for Health AI (CHAI). Accessed July 5, 2024. <https://chai.org/assurance-standards-guide/>
8. Microsoft Source. News Release. New consortium of healthcare leaders announces formation of Trustworthy & Responsible AI Network (TRAIN), making safe and fair AI accessible to every healthcare organization. Published March 11, 2024. Accessed April 18, 2024. <https://news.microsoft.com/2024/03/11/new-consortium-of-healthcare-leaders-announces-formation-of-trustworthy-responsible-ai-network-train-making-safe-and-fair-ai-accessible-to-every-healthcare-organization/>
9. Cary MP Jr, Zink A, Wei S, et al. Mitigating racial and ethnic bias and advancing health equity in clinical algorithms: a scoping review. *Health Aff (Millwood)*. 2023;42(10):1359-1368. doi: 10.1377/hlthaff.2023.00553
10. Gordon M, Daniel M, Ajiboye A, et al. A scoping review of artificial intelligence in medical education: BEME Guide No. 84. *Med Teach*. 2024;46(4):446-470. doi: 10.1080/0142159X.2024.2314198

Electronically published July 10, 2024.

Address correspondence to Michael J. Pencina, PhD, Duke University School of Medicine, 2424 Erwin Rd, Suite 903, Durham, NC 27705 (michael.pencina@duke.edu).

N C Med J. 2024;85(4):241-243. ©2024 by the North Carolina Institute of Medicine and The Duke Endowment. All rights reserved. 0029-2559/2024/85407