

Artificial intelligence and predictive algorithms in medicine

Promise and problems

Susan P. Phillips MD CCFP MSc Sheryl Spithoff MD CCFP MSc Amber Simpson PhD

Artificial intelligence (AI) is poised to reshape preventive medical practice; however, its benefits to patients, specific social groups (eg, racialized populations), and businesses remain to be seen. *Artificial intelligence* is defined as a “machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations, or decisions influencing real or virtual environments.”¹ Although AI does not play a major role in Canadian primary care yet, the merit of some algorithms developed elsewhere has been diminished by unintended and intended biases. For example, a US algorithm tasked to identify American patients whose current illnesses might predict future needs for increased care erroneously concluded that Black patients are healthier than equally sick White patients, effectively denying resources to Black patients.² The authors found that Black patients generated lower health care costs than White patients and that Black patients visited doctors less frequently than White patients; and yet, the AI was coded to interpret this less frequent access of care by Black patients as having a lesser disease burden.² Avoiding biases such as this requires awareness of risks and bold but thoughtful action by researchers and governments.

Promises

Artificial intelligence can account for the health effects of any available data set (eg, genetics, environment, behaviour); that is, for a patient’s context and biology, AI can “put the patient back together”³ to optimize and personalize prevention and diagnosis. It therefore has the potential to level the playing field and correct existing health inequalities.^{4,5} By incorporating volumes of real-world electronic medical record (EMR) data, AI addresses the effectiveness rather than only the efficacy of clinical trials, while concurrently incorporating key social determinants of health as predictors.⁶ By using EMRs of patients whose medical outcomes are known, algorithms can predict future risk of those outcomes in others.⁷ Artificial intelligence-generated algorithms could correct historic underrepresentation or misrepresentation of women and those of particular racial or ethnic backgrounds and could account for social circumstances like education, income, and social capital. For example, an algorithm developed to identify cardiovascular disease risk might include not only risk scores, but also socioeconomic status, which is a strong predictor of heart disease.⁸ Curating, aggregating, and analyzing longitudinal information derived from EMRs hold the promise of risk

prediction, prevention, and treatment by accounting for individual biology and context.

Problems

Artificial intelligence algorithms are not without medical, ethical, and social challenges, some because AI incorporates the values of data writers (eg, physicians), developers, and funders,⁹⁻¹¹ and others because entries in EMRs are neither standardized nor written to answer research questions. Information entered in EMRs is circumscribed by patient access to care (no clinical encounter means no recorded data about that patient) and clinicians’ values and determinations about that care. Ambiguity, lack of standard terminology, inaccuracy, and misleading acronyms are inevitable.¹² Electronic medical record data are limited by the time span of the record, and can be flawed for many reasons, including inconsistent measurement or documentation of biomarkers, absent modifiers of diagnosis such as severity, or even uninterpretable spelling. To merge known social and biological predictors of health, AI requires availability and precision of sociodemographic parameters such as race and socioeconomic status, data not routinely included in Canadian EMRs. This gap means algorithms must either ignore these data or risk imputing these classifications.¹³ Furthermore, recorded diagnoses, even those using ICD codes, are often vague (eg, does a breast cancer entry name the patient’s fear, the subject of a discussion, or a diagnosis?), as are terms such as *worse* and *improved*. Algorithms could use only biomedical measures without factoring in social determinants, but such oversight could diminish the “promise” of AI and might result in harmful blindness to the medical effects of social circumstances.¹⁴ Doctors likely consider a patient’s social determinants of health in addition to their health issues; failure to do so would, at most, harm 1 patient. Aggregation of EMR data to predict health risks aggregates blindness to sociodemographic characteristics and can undermine accurate predictions for many patients.^{2,15,16}

Misdiagnosis and overdiagnosis

Predictive AI models that are incomplete or unrepresentative can bias outcomes and precipitate misdiagnoses.¹⁷ Conversely, overdiagnosis may result from accepting a single abnormal test result as diagnostic, thereby attributing greater prognostic weight to some

measures (eg, blood glucose level in prediabetes) than is validated by research, narrowing definitions of what is normal (eg, lowering the upper limit of a normal hemoglobin A_{1c} level), or defining screening outcomes in ways that favour false-positive results over false-negative ones.^{9,18,19} All precipitate overdiagnosis and accompanying excessive investigation, patient fear, and unnecessary intervention. For example, if an algorithm's default is that "1 elevated blood pressure indicates risk," overdiagnosis will be inevitable, as more doctor visits will produce more recorded data, creating bias in what is, essentially, a time series with irregular and unequal sampling across participants.¹⁸

Artificial intelligence may, conversely, underdiagnose groups of people by excluding patients with missing data owing to limited access to care,² ignoring groups for whom data are not recorded or using biased disease definitions instead (eg, using a lower glomerular filtration rate to diagnose renal failure in Black versus White populations). Consider an algorithm to identify diabetes risk that selects as a marker of concern a blood glucose level recording that is higher than a specified level. As with current screening, those individuals who attend doctors infrequently will be tested less consistently and may be underdiagnosed. In terms of AI, if population subgroups do not attend physicians frequently and are underdiagnosed, their group characteristics (eg, age, race, sex) will be erroneously interpreted as conferring lower risk of the specific disease being tracked. This is 1 example of the more general possibility of distorting disease prediction when underlying data oversample those with specific medical traits, or are limited to either interesting or diagnosed cases rather than entire populations.²⁰

Who might have an interest in AI?

Although ethical principles of human rights, doing no harm, and protecting the individual are fundamentals of research, AI is only starting to be subjected to similar scrutiny.¹⁵ Without explicit patient consent, individuals whose medical records are electronic (almost 90% of primary care records in Canada in 2019²¹) have a digital health footprint that is being commodified, aggregated, shared, sold, and re-identified.²² Re-identification possibilities are magnified when EMR data are merged with self-quantification tools such as smart watches.²² Industry's drive to maximize business is a powerful underlying force in science and technology.²³ Marketers' self-interest can skew selection of input and output measures to amplify an advantageous (to industry) outcome by maximizing those identified as being at risk and by creating demand for unnecessary treatments. Statistical evidence has been misrepresented by looking at relative rather than absolute risk, again to inflate diagnoses and markets.^{5,24,25} Artificial intelligence algorithms that are proprietary, as most have been to date, tend to lack transparency, necessitating clever workarounds to audit them for such disparities

or distortion, and to find bias, inequality, and sources of misdiagnosis or overdiagnosis.^{2,24,26} The ethos of AI and precision medicine serves industry better than it does patients when it is propelled by and fosters the belief that "more is better," "new is better," "prevention is better than cure," and that diagnosing those at even minimal risk of disease is advantageous.

What is next?

Human bodies integrate social traits and life experiences with individual biology to interdependently determine health and illness.²⁵ Artificial intelligence algorithms are able to do this as well²⁴ but could exacerbate blindness to race, gender, and social identity, and obscure all behind a perceived certainty of science and numbers.⁴ An assessment is urgently needed of the empirical risks AI poses in terms of bias to inputs, processes, or the diagnoses that arise from outputs. The development of regulatory and practice frameworks to safeguard Canadians' health is also needed. This process is under way elsewhere and only beginning among Canadian researchers, physicians, ethicists, and regulators.²⁷⁻²⁹ Regulations and policies should address the commodification and monetization of health data; explicit patient consent; safeguarding of patient privacy; re-identification possibilities; the need to test algorithms using large, diverse, and representative data; transparent coding that enables auditing; when to inform patients that their diagnoses are AI generated; and chains of accountability for errors. Other jurisdictions are operationalizing oversight and regulations of both products (eg, algorithms) and processes (eg, data access, analytics) to ensure ethical standards and equitable resource allocation. Since 2010 in the United Kingdom, clinical decision support software requires approval from the Medicines and Healthcare Products Regulatory Agency. Developers must demonstrate that benefits outweigh risks, the effectiveness of the software, and adherence to criterion standards, and must conduct postmarket surveillance.³⁰ The US Food and Drug Administration has draft guidance for clinical decision support systems that use AI.^{31,32} Health Canada is considering regulatory requirements for AI as a medical device but not yet for stand-alone AI software systems.³³ A 2019 guide on the use of AI in clinical practice by the Canadian Medical Protective Association stated that responsibility to assess quality, functionality, reliability, and privacy of AI systems rests with the physician,³⁴ raising the question of whether ensuring efficacy, accountability, and privacy while not stymieing development of beneficial health frameworks should rest with individual doctors.

Conclusion

Only when there is assurance that values, AI, and underlying training and validation data sets align with and are as expansive and inclusive as are the bodies

they purport to help will their outputs foster social and racial equity and better health for all. The real question is whether regulators, data scientists, medical researchers, and clinicians will prioritize equity and ethics. 🌿

Dr Susan P. Phillips is Professor in the Department of Family Medicine and Director of the Centre for Studies in Primary Care at Queen's University in Kingston, Ont.

Dr Sheryl Spithoff is a family physician at Women's College Hospital in Toronto, Ont, and Assistant Professor in the Department of Family and Community Medicine at the University of Toronto. **Dr Amber Simpson** is a computer scientist at Queen's University and Canada Research Chair in Biomedical Computing and Informatics.

Competing interests

None declared

Correspondence

Dr Susan P. Phillips; e-mail phillip@queensu.ca

The opinions expressed in commentaries are those of the authors. Publication does not imply endorsement by the College of Family Physicians of Canada.

References

1. The technical landscape. In: *Artificial intelligence in society*. Paris, France: Organisation for Economic Co-operation and Development; 2019.
2. Obermeyer Z, Powers B, Vogeli C, Mullainathan S. Dissecting racial bias in an algorithm used to manage the health of populations. *Science* 2019;366(6464):447-53.
3. Greene JA, Loscalzo J. Putting the patient back together—social medicine, network medicine, and the limits of reductionism. *N Engl J Med* 2017;377(25):2493-9.
4. Geneviève LD, Martani A, Shaw D, Elger BS, Wangmo T. Structural racism in precision medicine: leaving no one behind. *BMC Med Ethics* 2020;21(1):17.
5. Chen IY, Joshi S, Ghassemi M. Treating health disparities with artificial intelligence. *Nat Med* 2020;26(1):16-7.
6. Phillips SP, Hamberg K. Doubly blind: a systematic review of gender in randomised controlled trials. *Glob Health Action* 2016;9:29597.
7. Suresh H, Guttag JV. A framework for understanding unintended consequences of machine learning. *ArXiv* 2019;1901.10002.
8. Marmot M, Shipley M, Brunner E, Hemingway H. Relative contribution of early life and adult socioeconomic factors to adult morbidity in the Whitehall II study. *J Epidemiol Community Health* 2001;55(5):301-7.
9. Carter SM, Rogers W, Win KT, Frazer H, Richards B, Houssami N. The ethical, legal and social implications of using artificial intelligence systems in breast cancer care. *Breast* 2020;49:25-32. Epub 2019 Oct 11.
10. Benjamin R. *Race after technology*. Cambridge, UK: Polity Press; 2019.
11. Noble SU. *Algorithms of oppression: how search engines reinforce racism*. New York, NY: New York University Press; 2018.
12. Chen IY, Szolovits P, Ghassemi M. Can AI help reduce disparities in general medical and mental health care? *AMA J Ethics* 2019;21(2):E167-79.
13. Gianfrancesco MA, Tamang S, Yazdany J, Schmajuk G. Potential biases in machine learning algorithms using electronic health record data. *JAMA Intern Med* 2018;178(11):1544-7.
14. Dubber MD, Pasquale F, Das S. *The Oxford handbook of ethics of AI*. Oxford, UK: Oxford University Press; 2020.
15. Tannenbaum C, Ellis RP, Eyssel F, Zou J, Schiebinger L. Sex and gender analysis improves science and engineering. *Nature* 2019;575(7781):137-46. Epub 2019 Nov 6.
16. Raji ID, Smart A, White RN, Mitchell M, Geburu T, Hutchinson B, et al. Closing the AI accountability gap: defining an end-to-end framework for internal algorithmic auditing. *ArXiv* 2020;2001.00973.
17. Vogt H, Green S, Ekstrøm CT, Brodersen J. How precision medicine and screening with big data could increase overdiagnosis. *BMJ* 2019;366:l5270.
18. Zheng K, Gao J, Ngiam KY, Ooi BC, Yip WLJ. Resolving the bias in electronic medical records. In: *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '17)*. New York, NY: Association for Computing Machinery; 2017.
19. Cahn A, Shoshan A, Sagiv T, Yesharim R, Goshen R, Shalev V, et al. Prediction of progression from pre-diabetes to diabetes: development and validation of a machine learning model. *Diabetes Metab Res Rev* 2020;36(2):e3252. Epub 2020 Jan 14.
20. Bae SH, Yoon KJ. Polyp detection via imbalanced learning and discriminative feature learning. *IEEE Trans Med Imaging* 2015;34(11):2379-93. Epub 2015 May 18.
21. *How Canada compares: results from the Commonwealth Fund's 2019 international health policy survey of primary care physicians*. Ottawa, ON: Canadian Institute for Health Information; 2020.
22. Grande D, Luna Marti X, Feuerstein-Simon R, Merchant RM, Asch DA, Lewson A, et al. Health policy and privacy challenges associated with digital technology. *JAMA Netw Open* 2020;3(7):e208285.
23. Sharon T. When digital health meets digital capitalism, how many common goods are at stake? *Big Data Soc* 2018;5(2):1-12.
24. Challen R, Denny J, Pitt M, Gompels L, Edwards T, Tsaneva-Atanasova K. Artificial intelligence, bias and clinical safety. *BMJ Qual Saf* 2019;28(3):231-7. Epub 2019 Jan 12.
25. Krieger N, Davey Smith G. "Bodies count," and body counts: social epidemiology and embodying inequality. *Epidemiol Rev* 2004;26:92-103.
26. Cirillo D, Catuara-Solarz S, Morey C, Guney E, Subirats L, Mellino S, et al. Sex and gender differences and biases in artificial intelligence for biomedicine and healthcare. *NPJ Digit Med* 2020;3:81.
27. Obermeyer Z, Nissam R, Stern M, Eaneff S, Bembeneck EJ, Mullainathan S. *Algorithmic bias playbook*. Chicago, IL: Chicago Booth Center for Applied Artificial Intelligence; 2021. Available from: <https://www.chicagobooth.edu/research/center-for-applied-artificial-intelligence/research/algorithmic-bias/playbook>. Accessed 2022 Jul 14.
28. Smith G, Rustagi I. *Mitigating bias in artificial intelligence: an equity fluent leadership playbook*. Berkeley, CA: Berkeley Haas School of Business; 2020. Available from: <https://haas.berkeley.edu/equity/industry/playbooks/mitigating-bias-in-ai/>. Accessed 2022 Jul 14.
29. Faes L, Liu X, Wagner SK, Fu DJ, Balaskas K, Sim DA, et al. A clinician's guide to artificial intelligence: how to critically appraise machine learning studies. *Transl Vis Sci Technol* 2020;9(2):7. Erratum in: *Transl Vis Sci Technol* 2020;9(9):33.
30. Vollmer S, Mateen BA, Bohner G, Király FJ, Ghani R, Jonsson P, et al. Machine learning and artificial intelligence research for patient benefit: 20 critical questions on transparency, replicability, ethics, and effectiveness. *BMJ* 2020;368:l6927. Erratum in: *BMJ* 2020;369:m1312.
31. *Artificial intelligence and machine learning in software as a medical device*. Silver Spring, MD: US Food and Drug Administration; 2021. Available from: <https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-and-machine-learning-software-medical-device>. Accessed 2022 Jul 11.
32. *Clinical decision support software: draft guidance for industry and Food and Drug Administration staff*. Silver Spring, MD: US Food and Drug Administration; 2019. Available from: <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/clinical-decision-support-software>. Accessed 2020 Jun 5.
33. Canadian Institutes of Health Research, Health Canada. *Introduction of artificial intelligence and machine learning in medical devices*. Ottawa, ON: Government of Canada; 2019. Available from: <https://cibr-irc.gc.ca/e/51459.html>. Accessed 2019 May 10.
34. *Can AI assist you with your clinical decisions? Looking at the benefits and risks of AI technologies in medicine*. Ottawa, ON: Canadian Medical Protective Association; 2019. Available from: <https://www.cmpa-acpm.ca/en/advice-publications/browse-articles/2019/can-ai-assist-you-with-your-clinical-decision>. Accessed 2022 Jul 11.

This article has been peer reviewed.

Can Fam Physician 2022;68:570-2. DOI: 10.46747/cfp.6808570

La traduction en français de cet article se trouve à <https://www.cfp.ca> dans la table des matières du numéro d'août 2022 à la page e230.