

## COMMENT OPEN



# What do medical students actually need to know about artificial intelligence?

Liam G. McCoy<sup>1,2</sup>, Sujay Nagaraj<sup>1,3</sup>, Felipe Morgado<sup>1,4</sup>, Vinyas Harish<sup>1,2</sup>, Sunit Das<sup>1,5</sup> and Leo Anthony Celi<sup>6,7,8</sup>

With emerging innovations in artificial intelligence (AI) poised to substantially impact medical practice, interest in training current and future physicians about the technology is growing. Alongside comes the question of what, precisely, should medical students be taught. While competencies for the clinical usage of AI are broadly similar to those for any other novel technology, there are qualitative differences of critical importance to concerns regarding explainability, health equity, and data security. Drawing on experiences at the University of Toronto Faculty of Medicine and MIT Critical Data's "datathons", the authors advocate for a dual-focused approach: combining robust data science-focused additions to baseline health research curricula and extracurricular programs to cultivate leadership in this space.

*npj Digital Medicine* (2020)3:86; <https://doi.org/10.1038/s41746-020-0294-7>

## INTRODUCTION

With emerging innovations in artificial intelligence (AI) poised to substantially impact medical practice, interest in training current and future physicians about AI is growing<sup>1</sup>. Alongside this interest comes the question of what, precisely, medical students should learn<sup>2</sup>. While competencies for the clinical usage of AI are broadly similar to those for any other novel technology in medicine, there are qualitative differences of critical importance to concerns regarding explainability, health equity, and data security<sup>3–5</sup>. We advocate for a dual-focused approach: combining robust, learner-centered AI additions to baseline curricula and extracurricular programs to cultivate leadership in this space.

## WHAT DO PHYSICIANS NEED TO UNDERSTAND ABOUT AI IN THE CLINICAL CONTEXT?

Most directly, physicians need to understand AI in the same way that they need to understand any technology impacting clinical decision-making. A physician utilizing MRI, for example, does not need to understand the particle spin physics differentiating T1 and T2 weighted scans, but they do need to be able to:

- (i) Use it—identify when the technology is appropriate for a given clinical context, and what inputs are required to receive meaningful results.
- (ii) Interpret it—understand and interpret the results with a reasonable degree of accuracy, including awareness of sources of error, bias, or clinical inapplicability.
- (iii) Explain it—be able to communicate the results and the processes underlying them in a way that others (e.g. allied health professionals and patients) can understand.

These skills take on particular nuances in the context of AI. For (i) and (ii), it is critical for physicians to appreciate the highly context-specific nature of AI, and the fact that performance in a

single restricted context may not always be transferable. It is also important to be aware of factors which may decrease the performance of algorithms for specific patient groups<sup>3</sup>.

AI has been commonly criticized for the "black box" effect—that is, the mechanism by which a model arrives at a decision may be indecipherable<sup>1</sup>. This lack of technical "explainability", however, does not discharge the obligations of (iii). To satisfy requirements of informed consent and clinical collaboration, a physician may be called upon to communicate their understanding of the origin, nature, and justification of an algorithm's results to patients, families, and colleagues.

## WHAT DO PHYSICIANS NEED TO UNDERSTAND ABOUT AI IN THE BROADER PROFESSIONAL CONTEXT?

The professional obligations of physicians extend beyond the clinical role into leadership and health advocacy. The disruptive prospects of AI in healthcare raise significant ethical and operational challenges which physicians must collectively be prepared to engage with for the sake of ensuring patient welfare.

Substantial concerns exist regarding the impact of algorithmic clinical decision support on health equity, due to factors such as the use of datasets lacking representation from minority populations<sup>3</sup>, and the possibility for algorithms to learn from and perpetuate existing biases<sup>4</sup>. Risks around data security and privacy are also becoming rapidly apparent<sup>5</sup>. There is also, however, the potential for AI itself to alleviate some of medicine's existing problems with bias and unfairness<sup>6</sup>. Physicians should be aware of both possibilities and be equipped to advocate for the development and deployment of ethical and equitable systems. Finally, physicians must act as responsible stewards for patient data to ensure that the foundational trust between provider and patient is not violated.

<sup>1</sup>Faculty of Medicine, University of Toronto, Medical Sciences Building, 1 King's College Cir, Toronto, ON M5S 1A8, Canada. <sup>2</sup>Institute of Health Policy, Management and Evaluation, Dalla Lana School of Public Health, University of Toronto, 155 College St 4th Floor, Toronto, ON M5T 3M6, Canada. <sup>3</sup>Department of Computer Science, University of Toronto, 40 St. George Street, Room 4283, Toronto, ON M5S 2E4, Canada. <sup>4</sup>Department of Medical Biophysics, University of Toronto, 101 College St, Suite 15-701, Toronto, ON M5G 1L7, Canada. <sup>5</sup>Centre for Ethics, University of Toronto, 15 Devonshire Pl, Toronto, ON M5S 1H8, Canada. <sup>6</sup>Institute for Medical Engineering and Science, Massachusetts Institute of Technology, 77 Massachusetts Avenue, E25-505, Cambridge, MA 02139, USA. <sup>7</sup>Division of Pulmonary, Critical Care and Sleep Medicine, Beth Israel Deaconess Medical Center, 330 Brookline Avenue, Boston, MA 02215, USA. <sup>8</sup>Department of Biostatistics, Harvard T.H. Chan School of Public Health, 677 Huntington Avenue, Boston, MA 02115, USA. ✉email: [liam.mccoy@mail.utoronto.ca](mailto:liam.mccoy@mail.utoronto.ca)

## HOW MIGHT MEDICAL STUDENTS LEARN WHAT THEY NEED TO LEARN?

Concerted efforts should be taken to cultivate physician-leaders who are fluent in both AI and medicine. Such dual competence is important, as it is no simple task to select clinically relevant and computationally feasible targets for AI in medicine. A siloed approach may lead to clear clinical targets going unnoticed and worsen the production of technical “solutions in search of problems”<sup>7</sup>. A multidisciplinary, integrated approach to learning will serve to facilitate this goal.

When approaching such a complex topic, it is critical to distinguish between that which all physicians must know for everyday practice, and that which some physicians should know to drive innovation. Curricular components should be targeted to address the former, while robust extracurricular programs can be targeted toward the latter. Both components serve to promote discussions on how the convergence between AI and medicine is currently impacting and will continue to impact the physician’s identity. This aligns with the concept of the “reimagined medical school”, which establishes a framework of core knowledge while supporting students who seek deep dives into specific subject areas<sup>8</sup>.

This approach has been piloted at the University of Toronto (UofT) Faculty of Medicine and has been embraced by administration as an important part of the Faculty’s strategic plan<sup>8</sup>. Lectures in the preclinical curriculum introduce all students to these concepts, and the 2-year-long “Computing for Medicine” certificate program

provides particularly interested students with practical programming skills and immersion into clinical data science projects<sup>9</sup>. Additionally, an “AI in Medicine” student interest group hosts extracurricular seminars on the subject and helps to facilitate connections between medical students and a city’s broader AI ecosystem (in academia and industry) (see Supplementary Table 1 for a list of AI in Medicine offerings in the last two years).

Harvard Medical School has engaged in a similar approach, offering clinical informatics training as an elective for medical students<sup>10</sup>. During this elective, students are paired with faculty mentors in their area of interest and engage in a mix of didactic and hands-on learning to explore how informatics is embedded into health systems. The School has also collaborated with the MIT Critical Data group to offer a project-based course on data science in medicine<sup>11</sup>. Extracurricularly, the MIT Critical Data Group has worked to spur interest in AI through “datathons” (brief competitions wherein computer scientists and clinicians work together to use data to solve clinical problems)<sup>12</sup>. These collaborations are emblematic of the possibilities for collaboration with non-medical faculties to enrich the education of medical students.

With insight from these experiences, we identify a series of important opportunities in both the curricular and extracurricular realms (outlined in Table 1). We wish to emphasize the importance of finding synergy between the learning objectives and their delivery, and of maintaining a learner-centered ethos with a focus upon student engagement rather than passive

**Table 1.** Potential curricular and extracurricular learning opportunities for artificial intelligence in medicine.

Curricular objective	Delivery recommendations	Extracurricular objective	Delivery recommendations
Promote physicians to be data-savvy consumers Students should be able to critically evaluate AI claims and understand the connection between models and clinical realities.	Actively engage students with hands-on workshops focused around: Recognizing appropriate potential applications of AI to health data Understanding how to discern between different methods that can be applied to data (e.g. the distinction between prediction and causal inference approaches)	Promote student interest groups in AI Interested students should be encouraged to connect and build networks around their shared AI focus.	Extend the broadly-used format of the “student interest group” to AI, enabling students to organize and autonomously host initiatives such as: Seminar events with prominent AI in Medicine speakers Hackathons and datathons in collaboration with computer science and engineering students
Instill durable fundamental concepts about AI, while avoiding technical specifics It is more important for students to have a robust conceptual understanding of AI and the structure of clinical data science than to understand constantly changing technical specifics.	Incorporate lecture and self-learning module content around: The basic pipeline of data acquisition, cleaning, analysis, and visualization Issues with data stewardship and data quality assurance in healthcare. Classes of machine learning approaches and common issues with design and integration of AI into clinical practice	Facilitate connections between medical students and industry in the health-AI space As AI in medicine is not siloed to solely academia or industry, students should have the opportunity to be exposed to the AI ecosystem in their local and broader communities	Leverage partnerships (either at the Faculty level or the student group level) to offer: Student site visits to start-ups to learn about entrepreneurship and the creation of health AI products and services Student research opportunities with health AI companies, or public-private partnerships
Introduce frameworks for approaching ethical considerations, both clinically and at a systems level Students should appreciate fairness, accountability, and transparency as core AI analogues to the traditional bioethics principles of beneficence, non-maleficence, autonomy, and justice <sup>15</sup> .	Students should participate in interactive case-based workshops and seminars lead by AI and ethics experts focused on: The special considerations AI requires at clinical and system levels in a case-based format How fairness, accountability, and transparency directly relate to core clinical values of beneficence, non-maleficence, autonomy, and justice which must permeate through all aspects of their care	Provide longitudinal programs to give students hands-on experience with real-world AI projects Theoretical knowledge should be supplemented with practical, real-world experience through formalized programs.	Longitudinal programs can include but are not limited to: “Computing for Medicine”, a validated 14-week course offered to preclinical students in UofT’s MD program to promote computer literacy, algorithmic thinking, and cross-domain collaboration <sup>9</sup> Non-technical projects involving AI or data science, such as using design-thinking approaches to implement existing AI tools into clinical practice and workflows
Promote computer science/data science as a dual-training path for MD/PhD and MD/MSc students Students should be provided with partnered, formalized learning opportunities that provide training at the intersection of health and data science.	Establish partnerships with institutes across computer science, biomedical engineering, the basic sciences, and public health, such as: U of T Faculty of Medicine’s partnerships with the Vector Institute for Artificial Intelligence, and Schwartz Reisman Institute for Technology and Society Harvard Medical School’s Collaborative Health Sciences and Technology MD/PhD Offerings with MIT	Encourage cross-disciplinary collaborations between medical students and data scientists Students should build interdisciplinary networks, and be encouraged to connect and collaborate with peers across faculties	Take active steps to break past disciplinary silos through initiatives such as: Shared “AI in Medicine” journal clubs open to both medical students and computer science/engineering students Collaborative events such as “datathons”, wherein ad-hoc interdisciplinary teams compete to answer clinical questions on open database

knowledge transfer. These concepts should be integrated with other aspects of the curriculum wherever appropriate (such as the inclusion of an AI case study in a workshop about ethical clinical decision-making), as the competencies required to effectively work with AI will often overlap with those required to fulfil other core aspects of the physician role such as advocacy, leadership, and communication. Medical schools have a critical role to play not only in helping their students learn but also in nurturing their academic interests and sowing the seeds of future leadership. These recommendations can and should be tailored to the context and strengths of each medical school, its partnerships, and its student body.

### WHAT ABOUT AFTER MEDICAL SCHOOL?

While detailed discussion on postgraduate medical education (PGME) and continuing medical education (CME) is outside the scope of this work, it is important to consider that medical education is viewed as a life-long pursuit and attention needs to be provided to learners at later career stages<sup>13</sup>. Competencies around AI could be integrated in PGME curricula in existing research or Quality Improvement (QI) blocks. Research training, for medical or surgical trainees, could be in technical areas such as data science or biomedical engineering but also in ethics, health services research, and medical education. QI would focus on translating and evaluating proven innovations into care. CME offerings through online or in-person workshops can not only allow clinicians to refresh their competencies over the course of their career but also empower established practitioners with the skills and knowledge to keep up with this field<sup>14</sup>. The various curricular aspects in Table 1 can be modified to suit learners at different stages in their careers.

### CONCLUSION

Ultimately, medical schools are tasked with training physicians for a future in which artificial intelligence is poised to play a significant role. In order to succeed at this task, it will be essential for students to have curricular and extracurricular learning opportunities around the clinical usage, technical limitations, and ethical implications of the tools at their disposal. Given the importance and potential impact of this technology, we must act both to ensure a base of artificial intelligence literacy among physicians at-large and to nurture the skills and interests of the future leaders who will drive innovation in this space.

Received: 26 January 2020; Accepted: 26 May 2020;

Published online: 19 June 2020

### REFERENCES

1. Topol, E. J. High-performance medicine: the convergence of human and artificial intelligence. *Nat. Med.* **25**, 44–56 (2019).
2. Wartman, S. A. The empirical challenge of 21st-century medical education. *Academic Med.* **94**, 1412–1415 (2019).
3. Adamson, A. S. & Smith, A. Machine learning and health care disparities in dermatology. *JAMA Dermatol* **154**, 1247–1248 (2018).
4. Parikh, R. B., Teeple, S. & Navathe, A. S. Addressing bias in artificial intelligence in health care. *JAMA*. <http://jamanetwork.com/journals/jama/fullarticle/2756196>. (2019)
5. Price, W. N. & Cohen, I. G. Privacy in the age of medical big data. *Nat. Med.* **25**, 37–43 (2019).
6. Chen, I. Y., Joshi, S. & Ghassemi, M. Treating health disparities with artificial intelligence. *Nat. Med.* **26**, 16–17 (2020).
7. Wiens, J. et al. Do no harm: a roadmap for responsible machine learning for health care. *Nat. Med.* **25**, 1337–1340 (2019).

8. Prober, C. G. & Khan, S. Medical education reimaged: a call to action. *Acad. Med.* **88**, 1407–1410 (2013).
9. Law, M., Veinot, P., Campbell, J., Craig, M. & Mylopoulos, M. Computing for medicine: can we prepare medical students for the future? *Acad. Med.* **94**, 353 (2019).
10. Harvard Medical School Course Catalogue. PD530.7 Clinical Informatics. [http://www.medcatalog.harvard.edu/coursedetails.aspx?cid=PD530.7&did=260&yid=2020&fbclid=IwAR3FRgDGVFK4ca\\_wHGgnXBwf3zRLkN8LmiJX8ph1q3tFc\\_g3-ZAVT5gK1qAI](http://www.medcatalog.harvard.edu/coursedetails.aspx?cid=PD530.7&did=260&yid=2020&fbclid=IwAR3FRgDGVFK4ca_wHGgnXBwf3zRLkN8LmiJX8ph1q3tFc_g3-ZAVT5gK1qAI) (2020).
11. MIT Critical Data. 2019.HST.953: Collaborative Data Science in Medicine. <https://criticaldata.mit.edu/blog/2019/08/06/hst-953-2019/>. (2020).
12. Aboab, J. et al. A “datathon” model to support cross-disciplinary collaboration. *Sci. Transl. Med.* **8**, 333ps8 (2016).
13. Aschenbrener, C. A., Ast, C. & Kirch, D. G. Graduate medical education: its role in achieving a true medical education continuum. *Acad. Med.* **90**, 1203–1209 (2015).
14. McMahon, G. T. The leadership case for investing in continuing professional development. *Acad. Med.* **92**, 1075–1077 (2017).
15. Floridi, L. et al. AI4People—an ethical framework for a good AI society: opportunities, risks, principles, and recommendations. *Minds Mach.* **28**, 689–707 (2018).

### ACKNOWLEDGEMENTS

We would like to acknowledge the Faculty of Medicine, MD Program, and Medical Society at the University of Toronto for their support and commitment to AI in Medicine Students’ Society and other initiatives driven by students in service of our profession and its changing needs. L.A.C. is funded by the National Institute of Health through the NIBIB R01 grant EB017205.

### AUTHOR CONTRIBUTIONS

Initial conceptions and design: L.G.M., S.N., V.H., F.M. L.A.C. Drafting of the paper: L.G.M., S.N., V.H., F.M. Critical revision of the paper for important intellectual content: L.G.M., S.N., V.H., F.M., S.D., L.A.C. All authors approve the submitted version of the paper and have agreed to be personally accountable for both their own contributions and the accuracy and integrity of the overall work.

### COMPETING INTERESTS

The authors declare no competing interests

### ADDITIONAL INFORMATION

**Supplementary information** is available for this paper at <https://doi.org/10.1038/s41746-020-0294-7>.

**Correspondence** and requests for materials should be addressed to L.G.M.

**Reprints and permission information** is available at <http://www.nature.com/reprints>

**Publisher’s note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2020