What Medical Students Should Know About Artificial Intelligence, From a Former Instructor

by Kali Konstantinopoulos, MS1

It's no secret that artificial intelligence (AI) in healthcare is a hot topic. AI has already found its way into multiple medical specialties, such as radiology and pathology, in which machine learning algorithms are able to identify possible diagnoses from image features. Now, with further advancements in the technology, physicians may be able to use AI to inform diagnoses based on patient vitals and lab results, predict which treatments will be most effective, and assist with documentation. Since AI has great potential to aid physicians in clinical decision-making and shift the future of medicine, the question of how best to incorporate AI into medical education arises: What do medical students need to know about AI? As a former teaching assistant for an introductory AI course at Indiana University (IU) Bloomington, this question is of great personal interest to me. Having reflected on my time teaching the subject and my experiences as a medical student, I believe that the most important thing medical students should know about AI is the foundational workflow of training and testing models, no matter how easy it is to get lost in the technical details.

International Business Machines (IBM) defines AI as "a field, which combines computer science and robust datasets, to enable problem-solving" (1). It can be thought of as a machine that can complete tasks that normally require human cognition, such as distinguishing objects in pictures or writing a piece of text. The subfield of AI that may be most relevant for medicine is machine learning, in which algorithms typically find patterns in provided data to classify items or identify subgroups within the dataset. The basic workflow of machine learning involves choosing and processing the data, using a subset of that data to train the model to find patterns, testing the algorithm on similar data it has not seen before, and then evaluating its performance. In medicine, a machine learning program may take in images of histological slides or radiological scans and use image features like color and shape to determine a potential pathology. In the future, we may use machine learning to help identify diagnoses, prognoses, and treatment options from physical exam and testing results.

When I was teaching machine learning in an engineering context, students' main concerns were usually related to technical details: preparing and analyzing the data, writing code, and optimizing the algorithm for improved accuracy. However, in a clinical setting, patients are front and center. In a 2020 article on AI medical education by McCoy et al., the authors reasoned that physicians should be able to "identify when [AI] is appropriate for a given clinical context...understand and interpret the results with a reasonable degree of accuracy, including awareness of sources of error, bias, or clinical inapplicability" and "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" (2). The authors advocate for an educational model in which AI topics necessary for everyday clinical practice are integrated into the medical school curriculum, while quantitative skills necessary to advance the technology are incorporated into extracurricular activities. There are some existing models for computing-related extracurricular and curricular activities in medical school. The University of Toronto, for example, created a Computing in Medicine certificate for medical students, which included programming classes, medical computing projects, and

seminars with experts in the field (3). At the IU School of Medicine, the Terre Haute campus offers an elective called Advanced Topics in Biomedical Informatics and Technology, which covers, among other topics, "principles of biomedical data representation, software design…mathematical modeling…[and] big data" (4). While not specific to AI, these options address the computing literacy that would be necessary for medical students interested in becoming leaders in the medical software space.

From all this discussion on AI in medical education, it seems easy to conclude that all medical students need lessons in coding. However, when I think about machine learning concepts that could be applied to clinical practice, I don't envision memorizing the mathematical structure of a model or optimizing an algorithm. Instead, I go back to the basic workflow of machine learning: data acquisition and processing, training, testing, and evaluating. Even when teaching undergraduates majoring in engineering, I always returned to this framework to focus on critical analysis of the software. From a clinical perspective, the workflow may be the most important concept that a physician could explain to a patient, as it covers how patient data will be used and assuages fears that the AI will replace the physician. After all, the algorithm cannot think for itself and makes decisions based on how it was trained. Thinking beyond the exam room, understanding these foundational concepts also allows physicians to better collaborate with technologists to meet clinical needs. Having a baseline understanding of the machine learning process aids in the cross-disciplinary communication necessary to establish clinical relevance and ensure that new technologies meet the needs of both physicians and patients. This can include evaluating sources of bias, as physicians can use their knowledge of social determinants of health and related disparities to recognize bias in the dataset or features used to train the model.

The world of AI can seem like a black box (and for some algorithms, that may be true). However, when AI comes into the clinic, knowledge of the field cannot be left solely to computer scientists and engineers. With a basic understanding of the machine learning workflow, future physicians can be poised to evaluate new algorithms for clinical use, explain basic concepts to patients, and communicate with technical experts. So, while diagnoses of algorithmic bugs are left to the programmers, medical students can still be well-equipped to face the changing landscape of healthcare. No coding experience required.

References:

(1) International Business Machines. What is Artificial Intelligence (AI)? [Internet]. IBM: [cited 2023 Mar 1]. Available from: https://www.ibm.com/topics/artificial-intelligence.

(2) McCoy LG, Nagaraj S, Morgado F, Harish V, Das S, Celi LA. What do medical students actually need to know about artificial intelligence?. NPJ Digital Medicine [Internet]. 2020 Jun 19 [cited 2023 Mar 1];3(1):86. Available from: https://www.nature.com/articles/s41746-020-0294-7.

(3) Law M, Veinot P, Campbell J, Craig M, Mylopoulos M. Computing for medicine: Can we prepare medical students for the future?. Academic Medicine [Internet]. 2019 Mar [cited 2023 Mar 1];94(3):353. Available from: https://journals.lww.com/academicmedicine/Fulltext/2019/03000/Computing_for_Medicine__Can_We_Prepare_Medical.26.aspx.

(4) Indiana University School of Medicine. 2022-2023 Electives. IUSOM: [cited 2023 Mar 1]. Available from: https://iusom.sitehost.iu.edu/electives/.