Machine learning:
broadening the scope of ethical questions

Machine Learning in Genomics

Virtual Workshop

Genomic Data Science Working Group

National Advisory Council for Human Genome Research

April 13 and 14, 2021

Pamela Sankar, PhD

Department of Medical Ethics and Health Policy

University of Pennsylvania

Ethical issues identified with healthcare machine learning

Issues embedded in ML technology

- Transparency
- Bias
- Explainability
- Accountability

Standard medical ethics concerns

- Privacy
- Fairness,
- Justice.
- Informed consent; autonomy

Mixed commentary

Machine learning: Is it the technology that will

- will transform healthcare delivery in the next decade. Reduce costs; improve quality. (Becker 2019)

Or, is it

- A giant with feet of clay? (Cabitza 2019)

Progress into clinical settings has been slow

- Examples of fully integrated machine learning models that drive clinical care **are rare.** Sendak. 2019 eGEMs

- [T]his development …allows clinicians with limited ML knowledge to apply ML models to their datasets, [still] **we were not successful in identifying any current applications in translational medicine.** Toh, Ebiomedicine 2020

- Given this impressive array of studies, it is perhaps surprising that **real world deployments of machine learning algorithms in clinical practice are rare.** Kelly BMC Medicine 2019

Unfamiliar with healthcare challenges

- [O]ne thing to ask [about an algorithm being developed] is, and it can be really scary to ask, is but like kind of the worst case scenario and I think that's a lot of the problem with healthcare is, you know, the worst case is that someone loses their life, you know, I mean or gets worse even, it's just… versus, you know, someone's Uber didn't show up. It's a different… it's a different worst case...[002]

- And so many of these AI groups just go into this big data like it's data and don't understand the consequences. I mean we're not… there's not much consequences in telling people to buy the wrong book….The… but there's a lot of consequences in telling people to do the wrong thing in healthcare…[055]

Not an acceptable ethos

I think there are some issues there, but I think the bigger… the disconnect that I've seen is in the ethos of how ... computer science has turned into a business**… has been built around this notion of tolerance to failure. The cost of doing it is cheap, so we can fail and move on, do another one. That is antithetical to medicine.** You're not allowed to fail with people. … [T]hat's not an acceptable ethos. [057]

Electronic health records

- And on the flip side of it, when they do send us labs and other clinical data, usually in extracts from an EMR system, then that half of the battle is trying to make sense of that really, really crazy data to… to make it fit alongside the algorithms that we've already built previously. [063]

- But they're interacting with like the user interface end of EPIC and they [physicians] … they'll say like well I just use the social history tab to code that, but then we're looking at massive database tables in the back end and trying to figure out where all that stuff they charted ended up. [049]

Start over?

[W]e need clinicians caring for patients, not entering data into systems, which for the most part do nothing with it… but this is a completely backwards approach, right. What we're basically doing, **instead of designing a system that says what do we need, let's go make sure we get… we have that data,** we're saying okay, why don't you just… why don't you doctor and nurse enter a ton of stuff, of which probably the majority is not very useful from an ML standpoint…[042]

…ad hoc design process, poor data hygiene, and a lack of statistical rigor

- A rigorous review of articles reporting studies of deep learning performance against health-care professionals in detecting diseases from medical imaging reported that of the 82 studies that qualified for detailed analysis **only four provided physicians with conditions that mirrored real clinical environments** (Liu et al Lancet 2019)

- An essay on pitfalls in machine learning research opened with the claim that research progress was hindered by reliance on "...an *ad hoc* design process, poor data hygiene, and a lack of statistical rigor in model evaluation." (Biderman NeurIPS 2020)

Factors slowing ML progress

Machine learning developers interviewed were

- unfamiliar with healthcare challenges;
- animated by a different ethos; and,
- bewildered by the complexity and multi-functionality of medical records

Machine learning as a general tool

Machine learning was created to “be a general tool, designed for general purposes in a general context.” (Gideon 2016)

- Machine learning emerged from efforts to create the capacity for computers to learn from their own operations using any data being processed from streams of information fed back by satellites to credit scores
- As it evolved, success was judged in part by how readily it could be adapted worked across settings.

Focus on tools obscures obstacles that slow progress

- ML developers who know very little about health care got so close to healthcare **because it never occurred them –or apparently anyone else—that they shouldn't or couldn't.**
- Moving from finance to healthcare would be no more difficult that it had been to move from finance to marketing.
- The hammer would find its nails.
- Data would be wrangled and algorithms trained.

Treat ML the source of ethical problems

- Thus, the machine learning algorithm might face difficulties at providing a reliable medical diagnosis for a patient coming from Eastern Asia….. (Grote 2019)

- Moreover, studies indicate that image recognition software is prone to develop biases (Grote 2019)

- If algorithms trained on data sets with these characteristics [older, white, males] are adopted in healthcare, they have the potential to exacerbate health disparities. (**Vayena 2018)**

- ….standard **machine learning** can acquire stereotyped biases from textual data (Caliskan, 2017)

Proposed solutions

- Most common solution proposed in these articles is to pay careful attention to samples used in test datasets
- What does this imply for attention to health disparities?

Integrating Ethics into Machine Learning for Precision Medicine

- University of Pennsylvania
  - Pamela Sankar (PI)
  - Sarah Evenson

- Stanford University
  - Mildred Cho (PI)
  - Meghan Halley
  - Ari Nichol
  - Carole Ann Frederico

- **Funded by NHGRI (R01 HGO 10476)**