What's in the Box? Issues in Algorithmic Medical Care

By Chuck Dinerstein — December 12, 2018

Courtesy Pixabay [1]

I've written [2] in the past about medicine by algorithm, but a recent article in Science Translational Medicine, as well as a deep dive into the writings of Daniel Dennett on the evolution of our minds [1], helped focus my thinking.

Regulatory Oversight

The FDA considers algorithmic medicine a form of medical device and approves their applications using the same standards. A new medical device is allowed to enter the market when clinical trials have demonstrated its safety and efficacy. But because clinical trials, by their nature, are controlled and involve highly controlled populations, these devices often perform differently when released into the larger community and the hands of many more physicians. Post-marketing surveillance is supposed to identify problems with devices after they have been released into the wild; our track record here is more spotty. Some devices and medications have been left in the marketplace for longer than warranted because alarms were missed or ignored resulting in substantial harm to patients, Essure, a form of birth control and Invokana, a medicine for diabetes come to mind. And to reduce regulatory burden, a new device may have its approval "fast-tracked" when it is similar to devices already approved. For example, changing design of stents, to improve their efficacy make use of similarity to gain earlier approval.

Algorithmic medicine's devices have two unique qualities making them different from the medical device we usually approve, the difficulty in understanding how they function, their opacity, and their ability to be quickly modified, their plasticity.

The Opacity Problem

The use of algorithms has taken two pathways in medicine. First, it can be used to strengthen and make more readily available skills we already have, the new ability of the Apple watch to detect atrial fibrillation is an example. In this instance, humans can readily interpret the why of the algorithm's decisions because the algorithm only improved or automated a model we provided. It is a top-down engineering problem, just like most medical devices on the market.

The second pathway is through the use of deep learning computer models, like IBM's Watson, of Jeopardy fame, or the systems that have beat humans at chess and Go. These are genuinely mysterious, opaque black boxes, we know that they work, but do not know how they arrive at their conclusions making explaining their behavior challenging. This is bottom-up design, we don't understand why it works, but it does and to understand its decisions we have to try and work backward from its conclusion.

This distinction between top-down and bottom-up design may be better characterized by asking which comes first comprehension or competence [3]. As it turns out, in the Darwinian world competence proceeds comprehension, it is a bottom-up world. Only in the past few hundred years has top-down design been able to extend our understanding to new competencies. The canonical example in medicine might be the discovery of penicillin when agar plates were ignored for a few days, and in a happenstance instance it was noted a colony of bacteria was surrounded by a zone with no life, antibacterial activity was present. This is often presented as a serendipitous moment transformed by the "prepared" mind, and it may well be. But it an example of competence without comprehension. We had no idea what the underlying mechanism of this antimicrobial competence was, but it was there to be seen. The next several years were filled with top-down, directed investigations to understand the underlying "science" and harness its power.

The Plasticity Problem

Algorithms "learn" from the underlying data we provide. Garbage in, garbage out is the term we might use, but to a large degree that is not the issue, the data sets have been "validated" and much of the garbage removed. But like the restricted population of clinical trials, the data used by the algorithm restricts its applicability. And algorithmic medicine, released into the wild may break down more quickly from edge cases, situations at the border or beyond the engineered uses of a device. [2] Post-marketing surveillance is meant to find those edge cases, but algorithms when fed more data, can change. Engineered devices are static; software is more dynamic, so the algorithm in the field can become much different than the algorithm that was initially approved.

What is the role of physicians?

Physicians currently provide knowledge and experience in the development of top-down algorithms and in structuring the data fed to deep learning systems. We serve a much smaller role in providing oversight of algorithmic diagnostic and therapeutic advice, a role that will grow.

> *"Providers help to balance an algorithm's goal of average, overall accuracy with the*

> *realities of individual patient experience and the possibilities of rare diseases outside algorithmic training."*

Physicians are the interface between competent but uncomprehending algorithms and patients. We can more readily identify the edge cases, where the algorithm breaks down; human judgment will make these systems more resilient. But to do this well requires improving and changing post-marketing surveillance. Algorithms should be tracked in real time, assessing their success and failures; but that requires greater transparency in development and reporting, which reduces if not eliminates "proprietary" algorithms and their financial benefits. We have made this mistake before, with electronic medical records, a mistake we continue to pay for financially and in the adverse effects on physicians and their relationship with patients.

> *"Insanity is doing the same thing over and over again and expecting a different result."*

[1] Dennett, a philosopher, has written extensively on the source of our consciousness and his From Bacteria to Bach and Back is one of his more accessible writings. I found that he put a great deal of neuroscience and both Western and Eastern philosophy into a great, thought-provoking read.

[2] The current regulation of edge cases involve Instructions for use (IFUs) that accompany all medical devices and indicate, often in great detail, the manner in which the device is to be used. For some devices, like surgical "robots," manufacturers require physicians to take a course and demonstrate competence before they are allowed to use them on their own. But physicians do not always follow IFUs, for example, the off-label use of drugs ignores the "IFU" for medications.

Source: Big Data and black-box medical algorithms Science Translational Medicine DOI:10.1126/scitranslmed.aao5333