Wachter sees at least one encouraging sign that progress is coming. In the past few years the behemoths of the tech world—Google, Amazon, Microsoft—have developed a strong interest in health care. Google, for example, partnered with researchers from U.C.S.F., Stanford University and the University of Chicago to develop models aimed at predicting events relevant to hospitalized patients, such as mortality and unexpected readmission.

To deal with the messy data problem, the researchers first translated data from two EHR systems into a standardized format called Fast Healthcare Interoperability Resources, or FHIR (pronounced "fire"). Then, rather than hand-selecting a set of variables such as blood pressure and heart rate, they had the model read patients' entire charts as they unfolded over time up until the point of hospitalization. The data unspooled into a total of 46,864,534,945 data points, including clinical notes. "What's interesting about that approach is every single prediction uses the exact same data to make the prediction," says Alvin Rajkomar, a physician and AI researcher at Google who led the effort. That element both simplifies data entry and enhances performance.

But the involvement of massive corporations also raises serious privacy concerns. In mid-November 2019 the *Wall Street Journal* reported that Google, through a partnership with Ascension, the country's second-largest health care system, had gained access to the records of tens of millions of people without their knowledge or consent. The company planned to use the data to develop machine-learning tools to make it easier for doctors to access patient data.

This type of data sharing is not unprecedented or illegal. Tariq Shaukat, Google Cloud's president of industry products and solutions, wrote that the data "cannot be used for any other purpose than for providing these services we're offering under the agreement, and patient data cannot and will not be combined with any Google consumer data." But those assurances did not stop the Department of Health and Human Services from opening an inquiry to determine whether Google/Ascension complied with Health Insurance Portability and Accountability Act regulations. As of press time, the inquiry was ongoing.

But privacy concerns should not halt the quest for better, smarter, more responsive electronic health records, according to Reider. There are ways to develop these systems that maintain privacy and security, he says.

Ultimately real transformation of medical practice may require an entirely new kind of EHR, one that is not simply a digital file folder. All the major EHRs are built on top of database-type architecture that is 20 to 30 years old, Reider observes. "It's rows and columns of information." He likens these systems to the software used to record inventory at a brick-and-mortar bookstore: "It would know which books it bought, and it would know which books it sold." Now envision how Amazon uses algorithms to predict what a customer might buy tomorrow and to anticipate demand. "They've engineered their systems so that they can learn in this way, and then they can autonomously take action," Reider says. Health care needs the same kind of transformative leap.

Cassandra Willyard is a science writer based in Madison, Wis.

Wiring Minds

Successfully applying AI to biomedicine requires innovators trained in contrasting cultures

By Amit Kaushal and Russ B. Altman

From the popular press to the largest health care conferences, promises of artificial intelligence revolutionizing biomedicine are ubiquitous. It often seems as if we are on the cusp of AI systems that can remotely identify a person about to get sick, make a diagnosis (no doctor needed!), select a custom AI-designed pharmaceutical and deliver it to the patient just in time—in an AI-powered self-driving car, of course.

If indeed this is the future, we are far from reaching it. To be sure, the pace of change has been rapid. Deep learning—the fast-growing subfield of AI that enables machines to diagnose pneumonia from chest x-rays or predict health deterioration from medical records—was unfamiliar even to most computer scientists a decade ago. And we do not know what evolutionary or revolutionary advances will drive AI in the coming decades. What we do know is that the success of biomedical AI depends not just on developing the technology but also on developing the people behind it.

Translating algorithmic advances to biomedical breakthroughs requires critically considering both realms of knowledge and endeavor on many levels. What, for example, are the true capabilities of a new technology, and what is simply hype? What problems in biomedicine are most likely to benefit from emerging computational capabilities? And how do we go from an interesting biomedical application of a new technology to the implementation of systems that actually improve human health? These challenging, multifaceted questions will need to be answered by interdisciplinary teams. The teams will require experts in AI, experts in biology and medicine, and, most important, leaders who can motivate and guide individuals with such diverse talents.

Unlike some domains in which AI has been applied, in biomedicine the consequences of failure are weighty. For a social media company, an AI model that is ineffective at increasing ad clicks can be detected and rolled back the same day. When it comes to medicine, however, human lives are at stake. Inadequately informed uses of AI can lead to obvious harm, such as inaccurate diagnostic or therapeutic recommendations, but also to more insidious failures, such as an algorithm that gives racially biased recommendations because it was trained with subtly biased data. Given the complexities of biomedicine and the inscrutable nature of many AI algorithms, it might be years before such a flaw is uncovered. Group leaders—whether in academia, pharmaceutical laboratories or start-ups—must not only understand the technical and scientific issues but also anticipate and articulate the potential risks, benefits and implications of the projects they undertake.

We need men and women who can build AI systems in med-

© 2019 Springer Nature Limited. All rights reserved.

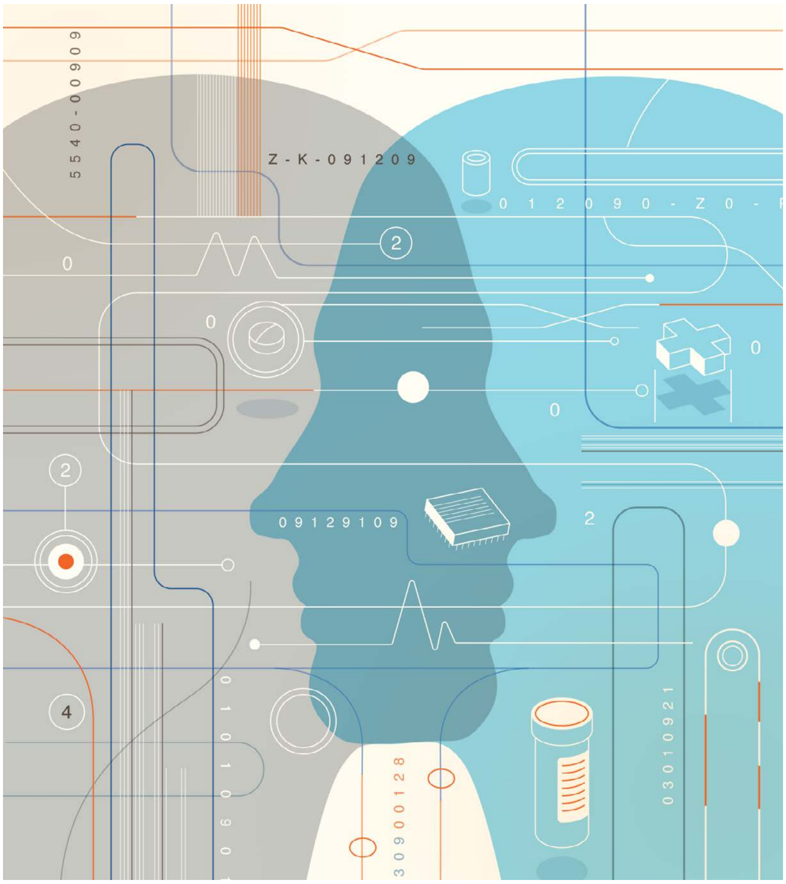

icine that improve care. It is relatively easy to generate excitement by solving the technical aspects of a problem, but making those advances useful often involves wrestling with the complex interplay of regulatory, economic and workflow issues in health care systems. Successful leaders benefit from deep knowledge and intuition in both the AI and the biomedical domains. But we face a critical shortage of such versatile individuals.

Tackling this gap is crucial to ensuring the long-term success of biomedical AI. A primary challenge is the length of study required in these disciplines, but a greater one is training students in two realms that could hardly be more different in their approaches to problem-solving. Computer science involves the quantitative rigor of mathematics, statistics and engineering, whereas biology is underpinned by the haphazard products of evolution. Properties of living things are, literally and figuratively, organic. We seek students with the intellectual flexibility and passion to undergo lengthy training in both these contrasting cultures. Are we asking for the impossible?

These individuals do exist, and their numbers are growing. The first approach to their training is to identify individuals who already have a deep background in either biomedical or computational science and then help them become skilled in the other area. Graduate programs (M.S., Ph.D. and M.D./Ph.D.) in biomedical informatics have filled this role since the early 1980s. These programs attract diverse students and have grown to include disciplines that go by various names: computational biology, bioinformatics, clinical informatics, biomedical data science, and so on. All are concerned with different applications of computer science to biomedicine.

But what about training students at the intersection of these disciplines even earlier in their careers—while their intellectual intuitions are still forming? The difference would be like that between learning a second language as an adult and growing up in a bilingual household: fluency is second nature for early starters.

In 2001 we launched an engineering major at Stanford University to enable undergraduates to learn computer science and statistics in the context of biology and medicine. The program creates graduates with a bachelor of science degree who have already wrestled intensively with the challenges of applying computational tools to hard problems in biomedicine. Our students take biology with premedical students and computer science with classmates who will work in Silicon Valley, and each completes a two- or three-quarter-long research project during his or her time at Stanford. They acquire knowledge with breadth across the biomedical and technical fields and depth in a narrower application area. At least one course on the societal and ethical implications of technology is also required.

After almost two decades of training biomedical-computation undergraduates, we can say that the model works. Many of our graduates have gone on to careers in academia, clinical medicine, start-up companies (both in and outside of the biology field), large companies, law firms, venture capital, and elsewhere. And the major has consistently drawn a 50–50 balance of men and women—true for only a minority of quantitatively intensive engineering majors.

For most, the major has shaped their professional identity: they are not "AI people doing bio" or "Bio people doing AI." Instead both of these intellectual traditions reside comfortably within their minds, each informing their understanding of the other. Whereas it is impossible to learn the entirety of biomedicine and computer science in just four years (or even in 40), these people move freely between the cultures of biology and computer science and have already learned to apply deep technical skills to the hardest societal challenges in biology and human health.

In addition to graduate programs, the development of a robust set of undergraduate programs at the interface of biomedicine and computation could give students who are in a formative period of their education the ability to move fluidly between these very different disciplines. Such programs would accelerate the emergence of the workforce required for appropriate use of AI to advance biology and health care.

Amit Kaushal is a clinical assistant professor of medicine and an adjunct professor of bioengineering at Stanford University.

Russ B. Altman is a professor of bioengineering, genetics, medicine and biomedical data science at Stanford University.

© 2019 Springer Nature Limited. All rights reserved.