



The role of education in AI (and vice versa)

Princeton's head of computer science, Jennifer Rexford, explains how artificial intelligence is set to change not only *what* teachers teach but *how* they teach.

The field of education is searching for ways to respond to the fast-approaching impact of artificial intelligence (AI) across all sectors and fields, including its own. In this conversation with Jennifer Rexford, computer-science chair at Princeton University, she explains how this requires not only focusing on teaching different skills to prepare today's youth for work in an AI world but also employing AI to adopt novel ways of teaching. She also reiterates a common call for including diversity of thought in developing AI to utilize it successfully. An edited transcript of Rexford's remarks follows.

Teaching skills for an AI future

There's tremendous anxiety about workforce issues in AI. And they build on larger, longer concerns about automation and the effect it will have on jobs.

What's different than [previous advances that have affected the workforce] is that machine learning and artificial intelligence are having a broad impact on a whole different set of fields. And, rather abruptly, this confluence of data, better computing power, and better algorithms happened so fast that the transformation is

moving quicker than governments, companies, and educational institutions can react and adapt.

If you flip the question around and ask, “What is AI *not* good at?”—in other words, “What are humans good at? What’s left for us to do?”—two things that are harder for the machine to do are exhibit creativity and social skills and perceptiveness.

For the roles that involve social perception—caregiving is a good example—you can automate some routine aspects, but a lot of other aspects really do rely on social skills and a perceptiveness that machines don’t have.

As for creativity, machines may behave in ways that *seem* creative, such as when they make a smart move in the game of Go or chess, but they’re really born of exhaustive enumeration and evaluation of the underlying data. It’s not born of that spark of creativity.

And that’s a clarion call for thinking about not only retraining but even basic education. The way we teach today, even at the K–6 or K–12 level, doesn’t put enough emphasis on creativity and social perceptiveness and design and working in teams. These are the kinds of things that are going to matter very much in the future and matter already, much faster than we’ve adapted to them.

Why diversity is vital to AI’s success

There’s certainly a significant underrepresentation problem in all areas of STEM [science, technology, engineering, and math], and this is higher in computer science than in some other areas and higher in artificial intelligence than in general computer science. So diversity is definitely a significant problem.

And there’s a tremendous shortage of skilled people to do this work. There’s this general cry of “We need more people, so we really can’t afford to leave talent on the table.”

But there’s a second piece to this: the technology being created is so important in so many lines of work and in so many businesses that leaving a large part of the population out of the conversation leaves these people ineffective in a whole bunch of areas in our society—not just computer science, but everywhere.

When on the plane to Davos, I read that there are automatic soap dispensers—the ones you stick your hands under—that don’t work for a lot of African-Americans. It turns out that many of these dispensers have cheap sensors that reflect light better off of lighter skin. And they were tested and evaluated on only a narrow portion of the demographic. That’s not conscious racism—someone didn’t necessarily pick a cheap sensor just to make it hard for African-Americans to get soap—but there obviously wasn’t a sufficient diversity of people involved in the design or the testing. You could imagine that wouldn’t have happened if 12 percent of the designers and 12 percent of the people involved in testing the equipment were African-American.

How Princeton works to attract greater diversity in AI

There are a few different things we can do to attract greater diversity, but none are a silver bullet. But one thing we do [at Princeton] is offer both an engineering and a Bachelor of Arts degree in computer science. Within the department, they’re identical, but the entry path is different: engineers take physics and chemistry in their freshman year, while our

BA students take a foreign language or other humanities courses in their freshman year. So these students don't have to know that they want to be an engineer when they walk in the door, and they don't have to take the prerequisites that frankly are less relevant to computer science than they are to some other engineering disciplines.

Our "late-entry path" to computer science through the BA is more than 50 percent female. That gives you an indication of how effective this program is at getting women into the degree.

And we've seen this program skyrocket in the past several years. What often happens is you take a computer-science course in your freshman year, you fall in love with it, you see its potential to influence other interests that you have, and you get hooked.

The second way we attract diversity is through our intro course, taken by 60 percent of all students, which is a very outward-facing course. Its focus is on interdisciplinary applications of computing and art, music, and science. The programs you write in the class are something you wouldn't be embarrassed to show your roommate or your parents. They're interesting. You turn your computer into a music keyboard. You use recursion, which is a concept in computer science to create art. I think that helps make computer science accessible. Forty or 50 percent of the students

in that course have never programmed before, so we have separate discussion sections that you can opt into if you've never coded before.

How AI can transform the way we teach

We're going to do an experiment next semester where we actually scan some students in the MRI machine—voluntarily, of course—while they watch flip lectures [video lectures], so that we can get a glimpse into how they learn, looking at both the differences among students and across time, and how that correlates with their understanding of the concepts and their concentration level while they're watching the videos. Learning *how* people learn will hopefully help us and others think more broadly about retraining down the road.

We can also apply machine learning to the neuroscience data we collect, or even to more mundane things such as logs of how many times students rewatch a portion of a video lecture or where they stumble or where they slow the lecture down to watch it more closely. Machine learning can help us get a better sense of exactly where they're struggling or where the material hasn't been explained well, or a mix of the two. They're related, obviously. And it can also potentially help us personalize the students' learning, so that they can learn at a more efficient pace than they can in today's one-size-fits-all classrooms. ♦

Jennifer Rexford was interviewed by **Rik Kirkland**, the senior managing editor of McKinsey Publishing, who is based in McKinsey's New York office.

Copyright © 2018 McKinsey & Company. All rights reserved.