



AI: The Fear, the Bias, and the Rewards

111 101 1111001
101 001 1011001
101 011 1011111

ROBOTICS HAVE HUGE POTENTIAL TO ENHANCE HEALTH, CULTURE, AND CIVILIZATION. BUT ONLY IF WE CLAMP DOWN ON ACCOUNTABILITY FOR MAKING SURE LEARNING MODELS ARE BUILT USING THE RIGHT DATA.

BY AYANNA HOWARD

Whether robots are friend or foe to humankind has long been debated. We have sci-fi writers and Hollywood to thank, in part, for both sides of the argument.

Indeed, there are some fairly disheartening dramas—think *Westworld*, *Blade Runner*, and *The Matrix*—that depict a dystopian future in which machines take over, leaving the joys and splendor of civilization in the dust.

At the same time, however, long-ago classics such as *The Jetsons*, the *Star Trek* TV series, and *2001: A Space Odyssey* inspired several generations of humans to dream big. After percolating for decades, many of those dreams are now becoming realities, and they're largely making the world a better place.

BEYOND SCIENCE FICTION

If not for the *Starship Enterprise's* holodeck, for example, would we have virtual and augmented reality helping us model and design buildings? Or providing a safe but realistic way to teach risky medical procedures? It's hard to say.

But there's no arguing that AI and robotics are enriching our world in countless ways: from allowing paraplegics to stand up and walk to helping reduce highway fatalities to exploring Mars as an alternative habitat for humankind.

What got me personally dreaming about what good technology could do was *The Bionic Woman*, a mid-70s TV show about an injured woman who gained extraordinary powers through artificial limbs. I decided then and there, between the ages of 11 and 12, that I wanted to create artificial limbs for people, too. I planned to go to medical school, but I soon discovered that biology—especially dissecting harmless, innocent frogs—wasn't my ideal cup of tea. Then I heard about engineering and realized that, if I became an engineer, I could do exactly what I wanted to do and build robots of the future, without the squeamish stomach.

Today, I design robots that help improve the lives of children in a variety of ways. But while I'm focused on how technology can make the world better, I also know that any new technology can be used not only for the noble but for the nefarious. A car can be a means to independence, employment, and leisure. But it can also be used as a killing machine.

HOW BIAS CREEPS IN

The ethics of what we create and how we develop it play a huge role in identifying what we call AI or robot bias. Those working towards a common good attempt to eliminate that bias the way we would in other aspects of

society involving human decision-making processes, such as in our hiring processes or in the criminal justice system. But we're not always successful. Bias can still sneak into the data we collect as we train deep-learning models, for example, or in the process of selecting attributes from the data for the algorithms to consider.

Stories of AI missing the mark abound. Some are comical, as in the case of wayward GPS systems. One recently caused a driver in Indonesia to drive off a cliff (he survived); in an older, scarier, incident, a bus in Washington State carrying a high school baseball team crashed into a wall while following GPS instructions. These mishaps occur when the data used in the modeling is the wrong data. For example, the data in these cases assumed a certain size vehicle when in fact the vehicles were larger and required a different route.

A prominent example of an AI project-gone-awry because of unintentional bias was an automated recruitment program that Amazon had to scrap. It was reported last fall that Amazon had been training its computer models to vet job applications based on a 10-year history of successful hires. But most applications from that 10-year period came from men, a reflection of the long-standing gender gap in the tech industry.

The data scientists on the project had simply overlooked this aspect of their data, and, as a result, Amazon's computer models were in effect being trained to filter out women. That wasn't the company's intent. So, it was back to the drawing board for Amazon.

AI AND ROBOTICS: TWO SIDES OF THE SAME COIN

AI and robotics are inextricably linked: I think of AI as the "brains" of the operation, while the physical robot is the "body" that carries out the work as instructed by the brain. And now that we have virtual robots—such as avatars in online worlds and disembodied digital assistants, such as Apple's Siri and Google's Alexa—the two disciplines have truly blended.

To work, as mentioned, AI and robotics need the right data—data that's been collected, prepared, and programmed in the context of a goal. And this is where different levels of AI bias can enter the picture.

We have to be careful that the objectives we set are ethical objectives. That the data we collect reflects reality (where the Amazon recruitment modeling broke down). It might not work, for example, to create a model based on data collected in California, apply it in Georgia, and expect it to work. And, finally, we need to be careful that the attributes the AI algorithms consider in calculating a

prediction or an action aren't biased, either in terms of the attributes we choose for the algorithm to consider or the attributes we choose to leave out.

PERSONALIZATION, PROFILING, AND PRIVACY

All these considerations dance around a fine line between using data for personalization—which most of us seem to want—and profiling, a word that now tends to make the hair on the back of people's necks stand up.

Nearly all the things we want AI to do for us today is about personalization, whether we're asking Siri to place a phone call, Alexa to play a song, or Waze to get us to our intended destination. Or we might be building an exoskeleton that will allow a paralyzed person to walk. Irrespective of the use case, getting the right result is all about programming the right data for the individual at hand.

And for successful programming, we have to make some generalizations. For example, if data about me says I'm a Georgia Tech professor on email till 10 p.m. every night, a marketing system might assume I'm a coffee drinker and target me accordingly. That could benefit me.

But when do we cross the line between "profiling" for everyone's benefit and making assumptions and generalizations that could offend, invade someone's privacy, or even do harm?

These are early days, and these lines of demarcation have not yet been fully drawn. But erroneous assumptions can result in AI bias, and that can have all kinds of unexpected consequences. In 2018 alone, for example, AI bias has caused immigrants to be erroneously deported, unsafe cancer treatments to be recommended, and an "ethnicity detection" feature to be created to search faces in New York City based on race without citizens' knowledge or permission.

These weren't AI's finest hours.

ABOUT THE AUTHOR:

Ayanna Howard is an American roboticist and Chair of the School of Interactive Computing at Georgia Institute of Technology. She's also the founder and CTO of Zyrobotics, LLC, a company that focuses on applying technology in ways that enhance the quality of life for children. Howard holds three patents, and among her many awards are: Brown Engineering Alumni Medal (BEAM), 2016; AAAS-Lemelson Invention Ambassador, 2016-2017; and Forbes' America's Top 50 Women In Tech, 2018.

Howard says her favorite robot of all time is Rosie, the frilly maid from "The Jetsons" cartoon.

PAYING FOR 'FREE' SERVICES WITH YOUR DATA

However, individuals have a role to play, too. If you're online and use a free service—whether its Facebook, Amazon, LinkedIn, Google, or some other digital entity—you've already given up much of your privacy. Most of us say we want our privacy, but are we willing to pay for it? For example, if Facebook started charging \$19.99 a month to guarantee user privacy, I don't know how many people would do it. We seem more willing to pay with our data, than with our dollars and cents.

As it stands today, we're basically giving these digital giants permission to use our information by clicking "I accept" at the bottom of a long list of terms and conditions. We can choose not to use any of these services, of course. We can choose not to read the news for free or call another country for free or find job opportunities for free. But are we willing to give up these capabilities to protect ourselves?

Fortunately, the level of consumer privacy responsibility assumed today by big digital companies is coming under scrutiny. It began with the General Data Protection Regulation (GDPR) last year in Europe. Now, in the U.S., the Social Media Privacy and Consumer Rights Act of 2018, proposed last spring, is under advisement at the U.S. Committee on Commerce, Science, and Transportation, which conducted its first hearing on the subject on Feb. 27 of this year. In addition, several states, including California, Washington, and Massachusetts, have been drafting their own privacy legislation in the absence of a federal privacy framework.

IT'S JUST THE DATA TALKING

In addition to their potential to invade your privacy, robots make people apprehensive because they think robots truly understand them. Rest assured, they don't. When a digital assistant or a robot takes an action that's eerily personal,

it's not because there's a machine reading your mind. It's because the right data has been entered into the computing model to return the desired result.

I work with robots that help children. A robot might be playing the role of exercise coach to a child who needs to move around more. There's more to it, though, than the robot performing the exercise that the child should emulate.

The robot needs to have characteristics that will get results the way an authority figure—such as a parent, gym teacher, or even a friend—would. That means that I need to create "emotional AI," whereby the kids want to please the robot, just as they would a friend or teacher.

As in real life, the robot exercise coach will become "frustrated" or "pleased" with the child based on his or her performance and stick-to-it-iveness. But though it might seem that the robot is doing all the thinking on its own, it's not. It's just the data talking.

AI'S EFFECT ON THE FUTURE OF WORK

There's no question that AI will change the future of work. Some changes are obvious: a dangerous or dirty job once performed by a human will be handled less precariously by a machine. Other changes are more subtle. If truck drivers are replaced by self-driving vehicles, there will be fewer fatalities on the highways caused by long-haul truckers falling asleep at the wheel. On the other hand, there are small cities along well-worn truck routes that survive on truckers stopping for gas, food, lodging, and other services. The fallout of that shift in the job landscape has yet to be calculated.

In the financial trading market, more AI researchers and data scientists are going to be hired than stockbrokers. If I'm trying to figure out what to invest in, there's so much more data now to consider, it's overwhelming. There are the traditional sources of data, but then there's public opinion via Twitter, information from Congressional hearings, and so on. It's impossible to consider it all without some AI modeling. We may still need a stockbroker to interact with the investor, but we'll need more AI and researchers at the back end.

And while the World Economic Forum anticipates that robots or some form of automation will conduct more than half of the work performed on planet earth by the year 2025, the organization also says that automation will create twice as many jobs as it eliminates. We have yet to fully envision where these jobs will come from.

DEMOCRATIZATION OF EDUCATION THROUGH AI?

I believe that as companies bring in new technologies, they'll also create the necessary training to develop worker skills and build a competent workforce. But the shifting job landscape will also affect traditional academia. We've come to a point in time where all college students should have at least a Computer Science (CS) 101 education.

Now, there's likely to be a shortage of instructors to teach CS 101. So, I believe we might be looking at the democratization of education through AI. In other words, I envision academic environments in which robots are called upon to teach CS 101 classes to the student masses.

Getting this education is important, because AI is here and it's going to have a big impact on our lives. In February, U.S. President Donald Trump signed an executive order asking federal government agencies to dedicate more resources and investment into research, promotion and training on AI. It's partly a defensive move: there could be negative intent by others using AI. If they're investing in it, and we're the good guys, we have to invest too. Otherwise, there will be an imbalance of power, weighted toward those who might want to use it against us.

Just as important is harnessing AI and robotics to advance humankind—whether that means robotics applied to improving human health, conquering space, or something else. Reaching our desired milestones requires that we exercise ethics and execute governance for AI accountability to help make sure we use the right data and prevent biases creeping into our models to avoid AI mishaps.

