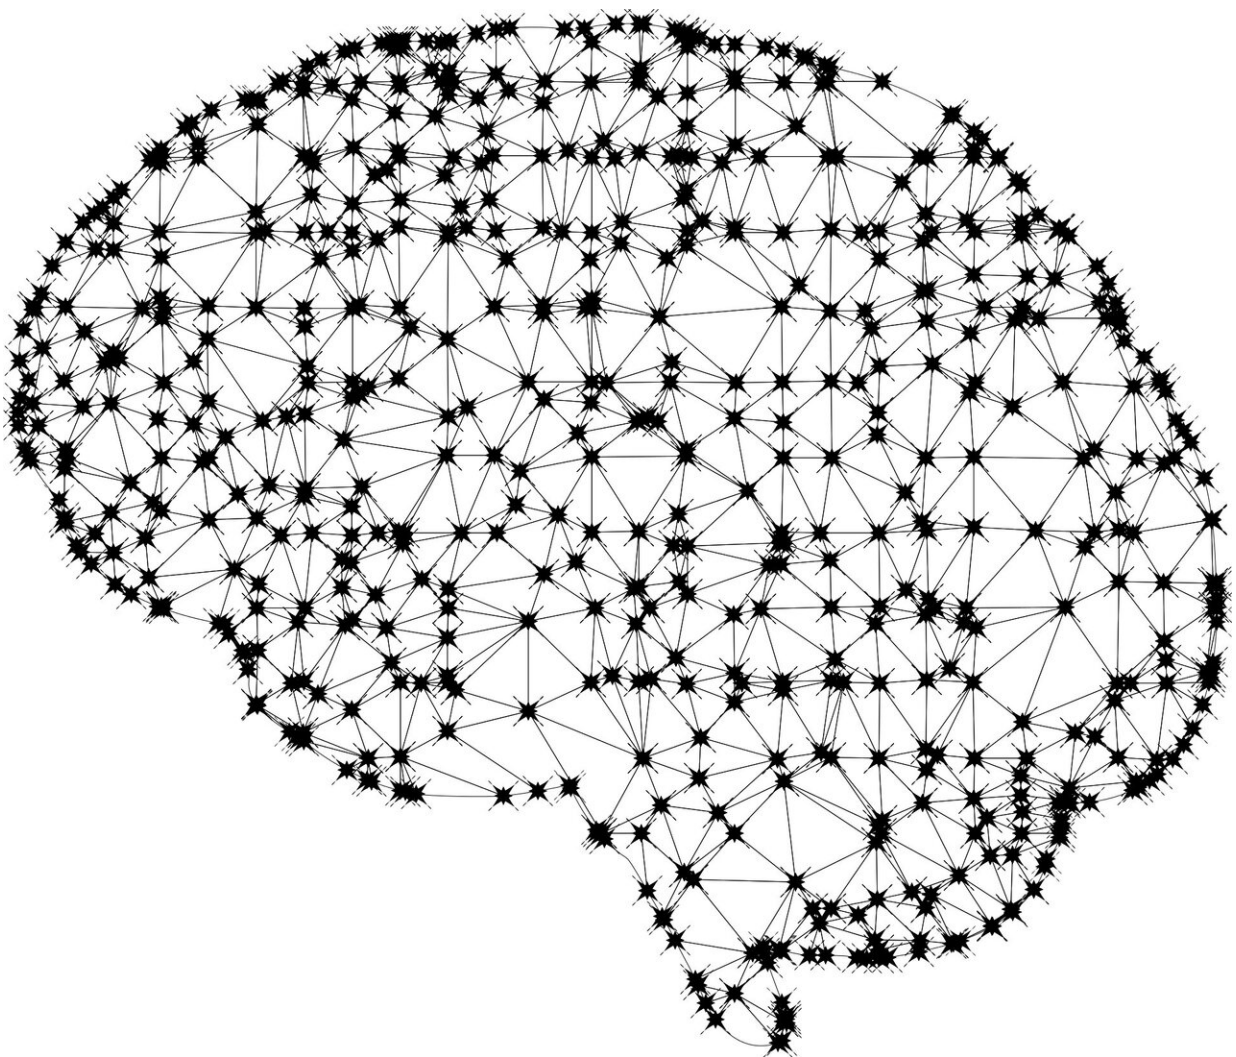


How computer algorithms help spread racial bias in US healthcare, and how they can help fix it

October 25 2019, by Amina Khan, Los Angeles Times



Credit: CC0 Public Domain

People may be biased, even without realizing it, but computer programs shouldn't have any reason to discriminate against black patients when predicting their health care needs. Right?

Wrong, new research suggests.

Scientists studying a widely used [algorithm](#) typical of the kind health insurers use to make crucial care decisions for millions of people have discovered significant evidence of racial bias when it comes to predicting the health risks of [black patients](#).

The findings, described Thursday in the journal *Science*, have far-reaching implications for the health and welfare of Americans as we become increasingly reliant on computers to turn raw data into useful information. The results also point to the root of the problem—and it isn't the computer program.

"We shouldn't be blaming the algorithm," said study leader Dr. Ziad Obermeyer, a machine learning and health researcher at the University of California, Berkeley. "We should be blaming ourselves, because the algorithm is just learning from the data we give it."

An algorithm is a set of instructions that describe how to perform a certain task. A recipe for brownies is an algorithm. So is the list of turns to make to drive to your friend's party.

A computer algorithm is no different, except that it's written in code instead of words. Today, they're used to target online ads, recognize faces and find patterns in large-scale data sets—hopefully turning the world into a more efficient, comprehensible place.

But while algorithms have become more powerful and ubiquitous, evidence has mounted that they reflect and even amplify real-world biases and racism.

An algorithm used to determine prison sentences was found to be racially biased, incorrectly predicting a higher recidivism risk for black defendants and a lower risk for white defendants. Facial recognition software has been shown to have both race and gender bias, accurately identifying a person's gender only among white men. Online advertisements that appear with Google search results have been found to show high-income jobs to men far more often than to women.

Obermeyer said it was almost by accident that he and his colleagues stumbled across the bias embedded in the health care algorithm they were studying.

The algorithm is used to identify patients with health conditions that are likely to lead to more serious complications and [higher costs](#) down the line. A large academic hospital had purchased it to help single out patients who were candidates for a care coordination program, which provides access to services such as expedited doctors' appointments and a team of nurses who may make house calls or refill prescriptions.

"It's kind of like a VIP program for people who really need extra help with their health," Obermeyer said.

The goal is to take care of these patients before their condition worsens. Not only does that keep them healthier in the long run, it keeps costs down for the health care system.

These kinds of algorithms are often proprietary, "making it difficult for independent researchers to dissect them," the study authors wrote. But in this case, the health system willingly provided it, along with data that

would allow researchers to see whether the algorithm was accurately predicting the patients' needs.

The researchers noticed something strange: Black patients that had been assigned the same high-risk score as white patients were far more likely to see their health deteriorate over the following year.

"At a given level of risk as seen by the algorithm, black patients ended up getting much sicker than white patients," Obermeyer said.

This didn't make sense, he said, so the scientists focused in on the discrepancy. They analyzed the health data from 6,079 black patients and 43,539 white patients and realized that the algorithm was doing exactly what it had been asked to do.

The problem was that the people who designed it had asked it to do the wrong thing.

The system evaluated patients based on the health costs they incurred, assuming that if their costs were high, it was because their needs were high. But the assumption that high costs were an indicator of high need turned out to be wrong, Obermeyer said, because black patients typically get fewer health care services than white patients, even when they're equally unwell.

That meant the algorithm was incorrectly steering some black patients away from the care coordination program.

Remedying that racial disparity could cause the percentage of black patients enrolled in the specialized care program to jump from 17.7% to 46.5%, the scientists realized.

Having identified the problem—a faulty human assumption—the

scientists set about fixing it. They developed one alternative model that zeroed in on "avoidable costs," such as emergency visits and hospitalizations. Another model focused on health, as measured by the number of flare-ups of chronic conditions over the year.

The researchers shared their discovery with the manufacturer of the algorithm, which then analyzed its national dataset of nearly 3.7 million commercially insured patients, confirming the results. Together, they experimented with a model that combined health prediction with cost prediction, ultimately reducing the bias by 84%.

Dr. Karen Joynt Maddox, a cardiologist and health policy researcher at Washington University of St. Louis, praised the work as "a thoughtful way to look at this really important emerging problem."

"We're increasingly putting a lot of trust on these algorithms and these black-box prediction models to tell us what to do, how to behave, how to treat patients, how to target interventions," said Joynt Maddox, who was not involved in the study. "It's unsettling, in a way, to think about whether or not these models that we just take for granted and are using are systematically disadvantaging particular groups."

The fault in this case was not with the algorithm itself, but with the assumptions made while designing it, she was quick to add.

Obermeyer said they chose not to single out the company that made the algorithm or the health system that used it. He said they hoped to emphasize the role of an entire group of risk-prediction algorithms that, by industry estimates, are used to evaluate roughly 200 million people a year.

Some people have reacted to discoveries of algorithmic bias by suggesting the algorithms be scrapped altogether—but the algorithms

aren't the problem, said Sendhil Mullainathan, a computational behavioral scientist at the University of Chicago and the study's senior author.

In fact, when properly studied and addressed, they can be part of the solution.

"They reflect the biases in the data that are our biases," Mullainathan said. "Now if you can figure out how to fix it ... the potential that it has to de-bias us is really strong."

A better algorithm may help to diagnose and treat the effects of racial disparities in care, but it cannot "cure" the disparity at the root of the problem: the fact that fewer dollars are spent on care of black patients, on average, than on white patients, he acknowledged.

These cost disparities likely arise for a number of reasons, the study authors said. Race and discrimination may play a role, and even when they have insurance, poorer patients face "substantial barriers to accessing health care."

"To the extent that race and socioeconomic status are correlated, these factors will differentially affect black patients," the researchers wrote.

Fixing those real-world sources of disparity presents a deeper and far more complicated challenge.

Ultimately, Obermeyer said, "it's a lot easier to fix bias in algorithms than in humans."

More information: Z. Obermeyer at University of California, Berkeley in Berkeley, CA et al., "Dissecting racial bias in an algorithm used to manage the health of populations," *Science* (2019).

[science.sciencemag.org/cgi/doi ... 1126/science.aax2342](https://science.sciencemag.org/cgi/doi/10.1126/science.aax2342)

©2019 Los Angeles Times

Distributed by Tribune Content Agency, LLC.

Citation: How computer algorithms help spread racial bias in US healthcare, and how they can help fix it (2019, October 25) retrieved 8 May 2024 from

<https://techxplore.com/news/2019-10-algorithms-racial-bias-healthcare.html>

<p>This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.</p>
--