



NEWS & COMMENTARY

Why Amazon's Automated Hiring Tool Discriminated Against Women



Rachel Goodman, Former Staff Attorney, ACLU Racial Justice Program

October 12, 2018

In 2014, a team of engineers at Amazon began working on a project to automate hiring at their company. Their task was to build an algorithm that could review resumes and determine which applicants Amazon should bring on board. But, according to a [Reuters](#) report this week, the project was canned just a year later, when it became clear that the tool systematically discriminated against women applying for technical jobs, such as software engineer positions.

It shouldn't surprise us at all that the tool developed this kind of bias. The existing pool of Amazon software engineers is overwhelmingly male, and the new software was fed data about those engineers' resumes. If you simply ask software to discover other resumes that look like the resumes in a "training" data set, reproducing the demographics of the existing workforce is virtually guaranteed.

In the case of the Amazon project, there were a few ways this happened. For example, the tool disadvantaged candidates who went to certain women's colleges presumably not attended by many existing Amazon engineers. It similarly downgraded resumes

that included the word “women’s” — as in “women’s rugby team.” And it privileged resumes with the kinds of verbs that men tend to use, like “executed” and “captured.”

Fortunately, Amazon stopped using the software program when it became clear the problem wasn’t going to go away despite programmers’ efforts to fix it. But recruiting tools that are likely similarly flawed are being used by hundreds of companies large and small, and their use is spreading.

[HOW ARTIFICIAL INTELLIGENCE IS CHANGING THE WORKPLACE](#)

There are many different models out there. Some machine learning programs — which learn how to complete a task based on the data they’re fed — scan resume text, while others analyze video interviews or performance on a game of some kind. Regardless, all such tools used for hiring measure success by looking for candidates who are in some way like a group of people (usually, current employees) designated as qualified or desirable *by a human*. As a result, these tools are not eliminating human bias — they are merely laundering it through software.

And it’s not just gender discrimination we should be concerned about. Think about all the ways in which looking at resume features might similarly cluster candidates by race: zip code, membership in a Black student union or a Latino professional association, or languages spoken. With video analysis, patterns of speech and eye contact have cultural components that can similarly lead to the exclusion of people from particular ethnic or racial groups. The same goes for certain physical or psychological disabilities.

We’ve seen these types of problems with artificial intelligence in [many other contexts](#). For example, when we used Amazon’s facial recognition tool to compare members of Congress against a database of mugshots, we got [28 incorrect matches](#) — and the rate for false matches was higher for members of color. This is due, in part, to the fact that the mugshot database itself had a disproportionately high number of people of color because of racial biases in the criminal justice system.

These tools are not eliminating human bias — they are merely laundering it through software.

Algorithms that disproportionately weed out job candidates of a particular gender, race, or religion are illegal under Title VII, the federal law prohibiting discrimination in employment. And that's true regardless of whether employers or toolmakers *intended* to discriminate — “disparate impact discrimination” is enough to make such practices illegal.

But it can be difficult to sue over disparate impact, particularly in “failure-to-hire” cases. Such lawsuits are very rare because it's so hard for someone who never got an interview to identify the policy or practice that led to her rejection.

That's why transparency around recruiting programs and other algorithms used by both companies and the government is so crucial. Many vendors who market these hiring tools claim that they test for bias and in fact are less biased than humans. But their software is proprietary, and there's currently no way to verify their claims. In some cases, careful work by outside auditors may be able to uncover bias, but their research is thwarted by various obstacles. We're [challenging](#) one such obstacle — a federal law that can criminalize testing of employment websites for discrimination.

But even this kind of outside research can't give us the full picture. We need regulators to examine not only the software itself but also applicant pools and hiring outcomes for companies that deploy the software. The Equal Employment Opportunity Commission, the federal agency that enforces laws against job discrimination, has begun to [explore](#) the implications of algorithms for fair employment, and we urge the agency to do more. EEOC should issue guidance for employers considering using these tools, detailing their potential liability for biased outcomes and steps they can take to test for and prevent bias. It should also include questions about data-driven bias in all of its investigations.

Big-data algorithms will replicate and even magnify the biases that exist in society at large — unless they are designed and monitored very, very carefully. The right kind of

oversight is required to make sure that happens.

Learn More About the Issues on This Page

Women's Rights

Women's Rights in the Workplace

Privacy & Technology

Related Content

Press Release



At Supreme Court, Tennessee Seeks to Expand Reach of Dobbs to Ban Health Care for Transgender Youth

Press Release



Civil Rights Groups Tell Federal Appeals Court that Protections for Pregnant Workers Cover Abortion Care

Podcast