

The sources of algorithmic bias

Description

“The foundations of algorithmic bias”, by Zachary Chase Lipton, 2016. pdf copy [here](#).
Original source [here](#)

[Tweet](#)

This morning, millions of people woke up and impulsively checked Facebook. They were greeted immediately by content curated by Facebook’s newsfeed algorithms. To some degree, this news might have influenced their perceptions of the day’s news, the economy’s outlook, and the state of the election. Every year, millions of people apply for jobs. Increasingly, their success might lie, in part, in the hands of computer programs tasked with matching applications to job openings. And every year, roughly 12 million people are arrested. Throughout the criminal justice system, computer-generated risk-assessments are used to determine which arrestees should be set free. In all these situations, algorithms are tasked with making decisions.

Algorithmic decision-making mediates more and more of our interactions, influencing our social experiences, the news we see, our finances, and our career opportunities. We task computer programs with approving lines of credit, curating news, and filtering job applicants. Courts even deploy computerized algorithms to predict “risk of recidivism”, the probability that an individual relapses into criminal behavior. It seems likely that this trend will only accelerate as breakthroughs in artificial intelligence rapidly broadened the capabilities of software.

Turning decision-making over to algorithms naturally raises worries about our ability to assess and enforce the neutrality of these new decision makers. How can we be sure that the algorithmically curated news doesn’t have a political party bias or job listings don’t reflect a gender or racial bias? What other biases might our automated processes be exhibiting that that we wouldn’t even know to look for?

Rick Davies Comment: This paper is well worth reading. It starts by explaining the basics (what an algorithm is and what machine learning is). Then it goes into detail about three sources of bias: (a) biased data, (c) bias by omission, and (c) surrogate objectives. It does not throw the baby out with the bathwater, i.e. condemn the use of algorithms altogether because of some bad practices and weaknesses in their use and design

[TAKEAWAYS]

Many of the problems with bias in algorithms are similar to problems with bias in humans. Some articles suggest that we can detect our own biases and therefore correct for them, while for machine learning we cannot. But this seems far fetched. We have little idea how the brain works. And ample studies show that humans are flagrantly biased in college admissions, employment decisions, dating behavior, and more. Moreover, we typically detect biases in human behavior post hoc by evaluating human behavior, not through an a priori examination of the processes by which we think.

Perhaps the most salient difference between human and algorithmic bias may be that with human decisions, we expect bias. Take for example, the well documented racial biases among employers, less likely to call back workers with more more typically black names than those with white names but identical resumes. We detect these biases because we suspect that they exist and have decided that they are undesirable, and therefore vigilantly test for their existence.

As algorithmic decision making slowly moves from simple rule based systems towards more complex, human level decision making, it's only reasonable to expect that these decisions are susceptible to bias.

Perhaps, by treating this bias as a property of the decision itself and not focusing overly on the algorithm that made it, we can bring to bear the same tools and institutions that have helped to strengthen ethics and equality in the workplace, college admissions etc. over the past century.

See also:

- [How to Hold Algorithms Accountable](#), Nicholas Diakopoulos and Sorelle Friedler. MIT Technology Review, November 17, 2016. Algorithmic systems have a way of making mistakes or leading to undesired consequences. Here are five principles to help technologists deal with that.
- [Is the Gig Economy Rigged?](#) by Will Knight November 17, 2016 A new study suggests that racial and gender bias affect the freelancing websites TaskRabbit and Fiverr and may be baked into underlying algorithms.

Category

1. Uncategorized

Date

30/04/2026

Date Created

20/11/2016

Author

admin