

Facebook's ad delivery system still discriminates by race, gender, age

December 19 2019, by Molly Callahan



Credit: Hannah Moore/Northeastern University

In settling five different lawsuits at once, Facebook earlier this year promised to [change the way it manages](#) the advertisements for housing, employment, and credit that run on its platform. Advertisers would no

longer be able to target an audience by gender, race, or age, in order to prevent against discrimination of legally protected groups.

But, [new research](#) by a team of computer scientists, including Northeastern professor Alan Mislove, shows that the algorithm Facebook uses to deliver advertisements can still skew toward specific demographic groups—despite the changes the company made.

"To be clear, we're not accusing Facebook of being in violation of the settlement," says Mislove, who is a professor of computer science at Northeastern. "Rather, what our research shows is a complex algorithm at work."

That's because algorithms, Mislove says, act much differently in practice than they do in theory.

In theory, removing the ability of advertisers to specifically target people by race, gender, and age should result in advertising audiences that include a diverse mix of people. In practice, Facebook's algorithm relies on myriad other characteristics about its users that ultimately serve as proxies for race, gender, and age, Mislove says.

This summer, Facebook introduced its Special Ad Audience feature—a [marketing tool](#) for people promoting ads that offer credit, employment, or housing opportunities. Whereas a clothing company might want to advertise directly to men or women, the domains of credit, employment, and housing have special legal protection in the United States to prevent discrimination.

Mislove worked with a team of researchers that also included Northeastern doctoral candidates Piotr Sapiezynski and Avijit Ghosh, undergraduate student Levi Kaplan, and a researcher from the nonprofit technology organization Upturn to test the new feature against its

predecessor, called a [Lookalike Audience](#)>.

The researchers used publicly available voter data to create audiences that were intentionally biased by race, gender, age, and [political views](#), and fed them to both the new and existing advertising tools to test whether the corresponding algorithm would reproduce each bias.

The results were striking: both Lookalike and Special Ad audiences replicated the demographic skews.

In a [statement to ProPublica](#), Facebook spokesman Joe Osborne said, "We've gone above and beyond others to help prevent discrimination in ads by restricting targeting and adding transparency. An advertiser determined to discriminate against people can do so on any online or offline medium today, which is why laws exist... We are the only digital media platform to make such meaningful changes in ads and we're proud of our progress."

Mislove says the results illustrate the difficult task of ensuring fairness in algorithms.

"In practice, an [algorithm](#) is given millions of inputs, each one of which is correlated in other ways to these protected features," Mislove says.

Information such as where you went to college, where you live, which pages you've 'liked' on Facebook, and more can inadvertently serve as indicators about your race, gender, and age—and algorithms that are built to optimize results may rely upon such information to determine who is most likely to click on an advertisement.

"Algorithms don't care," Mislove says. "They have a specific objective, and they're going to use the combination of features that will result in completing that objective; it doesn't matter to them if that means they're

delivering ads for tech jobs only to white men between the ages of 18 and 24."

This is a complex, but pressing problem, Mislove says, as algorithms underpin more and more of our daily activities. Algorithms are used to help make decisions about healthcare, credit worthiness, employment, criminal sentencing, travel directions, which music you listen to, and a dizzying array of other daily interactions.

All of the players in these domains "want the efficiencies of algorithms, but not everyone is considering the ways that they could go wrong," Mislove says.

When things do go wrong, the solution is often very complex.

"It's very hard right now, in the sense that the protected identities [against which discrimination can occur] really permeate our society," Mislove says. "It's going to be much harder and much more subtle than simply removing certain features at the outset."

More information: Algorithms That "Don't See Color": Comparing Biases in Lookalike and Special Ad Audiences:
[sapiezynski.com/papers/sapiezy ... ki2019algorithms.pdf](https://sapiezynski.com/papers/sapiezy...ki2019algorithms.pdf)

Provided by Northeastern University

Citation: Facebook's ad delivery system still discriminates by race, gender, age (2019, December 19) retrieved 24 April 2024 from <https://techxplore.com/news/2019-12-facebook-ad-delivery-discriminates-gender.html>

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.