


# Facial recognition software is biased towards white men, researcher finds



New research out of MIT's Media Lab is underscoring what other experts have reported or at least suspected before: facial recognition technology is subject to biases based on the data sets provided and the conditions in which algorithms are created.

Joy Buolamwini, a researcher at the MIT Media Lab, recently built a dataset of 1,270 faces, using the faces of politicians, selected based on their country's rankings for gender parity (in other words, having more females in office). Buolamwini then tested the accuracy of three facial recognition systems: those made by Microsoft, IBM, and Megvii of China. The results, which were originally reported in *The New York Times*, showed inaccuracies in gender identification dependent on a person's skin color.

Gender was misidentified in less than one percent of lighter-skinned males; in up to seven percent of lighter-skinned

females; up to 12 percent of darker-skinned males; and up to 35 percent in darker-skinner females.

“Overall, male subjects were more accurately classified than female subjects replicating previous findings (Ngan et al., 2015), and lighter subjects were more accurately classified than darker individuals,” Buolamwini wrote in a paper about her findings, which was co-authored by Timnit Gebru, a Microsoft researcher. “An intersectional breakdown reveals that all classifiers performed worst on darker female subjects.”

It’s hardly the first time that facial recognition technology has been proven inaccurate, but more and more evidence points towards the need for diverse data sets, as well as diversity among the people who create and deploy these technologies, in order for the algorithms to accurately recognize individuals regardless of race or other identifiers.

Back in 2015, Google was called out by a software engineer for mistakenly identifying his black friends as “gorillas” in its Photos app, something the company promised to fix (when in reality, it may have just removed the word “gorillas” from its index of search results in the app).

Two years ago, *The Atlantic* reported on how facial recognition technology used for law enforcement purposes may “disproportionately implicate African Americans.” It’s one of the larger concerns around this still-emerging technology – that innocent people could become suspects in crimes because of inaccurate technology – and something that Buolamwini and Gebru also cover in their paper, citing a year-long investigation across 100 police departments that revealed “African-American individuals are more likely to be stopped by law enforcement and be subjected to face recognition searches than individuals of other ethnicities.”

And, as *The Atlantic* story points out, other groups have found

in the past that facial recognition algorithms developed in Asia were more likely to accurately identify Asian people than white people; while algorithms developed in parts of Europe and the US were able to identify white faces better.

The algorithms aren't *intentionally* biased, but more research supports the notion that a lot more work needs to be done to limit these biases. "Since computer vision technology is being utilized in high-stakes sectors such as healthcare and law enforcement, more work needs to be done in benchmarking vision algorithms for various demographic and phenotypic groups," Buolamwini wrote.

<https://www.theverge.com/2018/2/11/17001218/facial-recognition-software-accuracy-technology-mit-white-men-black-women-error>