

Face recognition researcher fights Amazon over biased AI

April 3 2019, by Matt O'brien



In this Wednesday, Feb. 13, 2019, photo, Massachusetts Institute of Technology facial recognition researcher Joy Buolamwini takes questions from reporters at the school, in Cambridge, Mass. Buolamwini's research has uncovered racial and gender bias in facial analysis tools sold by companies such as Amazon that have a hard time recognizing certain faces, especially darker-skinned women. (AP Photo/Steven Senne)

Facial recognition technology was already seeping into everyday

life—from your photos on Facebook to police scans of mugshots—when Joy Buolamwini noticed a serious glitch: Some of the software couldn't detect dark-skinned faces like hers.

That revelation sparked the Massachusetts Institute of Technology researcher to launch a project that's having an outsized influence on the debate over how artificial intelligence should be deployed in the real world.

Her tests on software created by brand-name tech firms such as Amazon uncovered much higher error rates in classifying the gender of darker-skinned women than for lighter-skinned men.

Along the way, Buolamwini has spurred Microsoft and IBM to improve their systems and irked Amazon, which publicly attacked her research methods. On Wednesday, a group of AI scholars, including a winner of computer science's top prize, launched a spirited defense of her work and called on Amazon to stop selling its facial recognition software to police.

Her work has also caught the attention of political leaders in statehouses and Congress and led some to seek limits on the use of computer vision tools to analyze human faces.

"There needs to be a choice," said Buolamwini, a graduate student and researcher at MIT's Media Lab. "Right now, what's happening is these technologies are being deployed widely without oversight, oftentimes covertly, so that by the time we wake up, it's almost too late."



In this Wednesday, Feb. 13, 2019, photo, Massachusetts Institute of Technology facial recognition researcher Joy Buolamwini stands for a portrait at the school, in Cambridge, Mass. Buolamwini's research has uncovered racial and gender bias in facial analysis tools sold by companies such as Amazon that have a hard time recognizing certain faces, especially darker-skinned women. Buolamwini holds a white mask she had to use so that software could detect her face. (AP Photo/Steven Senne)

Buolamwini is hardly alone in expressing caution about the fast-moving adoption of facial recognition by police, government agencies and businesses from stores to apartment complexes. Many other researchers have shown how AI systems, which look for patterns in huge troves of data, will mimic the institutional biases embedded in the data they are learning from. For instance, if AI systems are developed using images of mostly white men, the systems will work best in recognizing white men.

Those disparities can sometimes be a matter of life or death: One recent study of the computer vision systems that enable self-driving cars to "see" the road shows they have a harder time detecting pedestrians with darker skin tones.

What's struck a chord about Boulamwini's work is her method of testing the systems created by well-known companies. She applies such systems to a skin-tone scale used by dermatologists, then names and shames those that show racial and gender bias. Buolamwini, who's also founded a coalition of scholars, activists and others called the Algorithmic Justice League, has blended her scholarly investigations with activism.

"It adds to a growing body of evidence that facial recognition affects different groups differently," said Shankar Narayan, of the American Civil Liberties Union of Washington state, where the group has sought restrictions on the technology. "Joy's work has been part of building that awareness."

Amazon, whose CEO, Jeff Bezos, she emailed directly last summer, has responded by aggressively taking aim at her research methods.

A Buolamwini-led study published just over a year ago found disparities in how facial-analysis systems built by IBM, Microsoft and the Chinese company Face Plus Plus classified people by gender. Darker-skinned women were the most misclassified group, with error rates of up to 34.7%. By contrast, the maximum error rate for lighter-skinned males was less than 1%.



In this Wednesday, Feb. 13, 2019, photo, Massachusetts Institute of Technology facial recognition researcher Joy Buolamwini stands for a portrait at the school, in Cambridge, Mass. Buolamwini's research has uncovered racial and gender bias in facial analysis tools sold by companies such as Amazon that have a hard time recognizing certain faces, especially darker-skinned women. Buolamwini holds a white mask she had to use so that software could detect her face. (AP Photo/Steven Senne)

The study called for "urgent attention" to address the bias.

"I responded pretty much right away," said Ruchir Puri, chief scientist of IBM Research, describing an email he received from Buolamwini last year.

Since then, he said, "it's been a very fruitful relationship" that informed IBM's unveiling this year of a new 1 million-image database for better analyzing the diversity of human faces. Previous systems have been

overly reliant on what Buolamwini calls "pale male" image repositories.

Microsoft, which had the lowest error rates, declined comment. Messages left with Megvii, which owns Face Plus Plus, weren't immediately returned.

Months after her first study, when Buolamwini worked with University of Toronto researcher Inioluwa Deborah Raji on a follow-up test, all three companies showed major improvements.

But this time they also added Amazon, which has sold the system it calls Rekognition to law enforcement agencies. The results, published in late January, showed Amazon badly misidentifying darker-hued women.



In this Wednesday, Feb. 13, 2019, photo, Massachusetts Institute of Technology facial recognition researcher Joy Buolamwini stands for a portrait behind a mask at the school, in Cambridge, Mass. Buolamwini's research has uncovered racial and gender bias in facial analysis tools sold by companies such as Amazon that

have a hard time recognizing certain faces, especially darker-skinned women. Buolamwini holds a white mask she had to use so that software could detect her face. (AP Photo/Steven Senne)

"We were surprised to see that Amazon was where their competitors were a year ago," Buolamwini said.

Amazon dismissed what it called Buolamwini's "erroneous claims" and said the study confused facial analysis with facial recognition, improperly measuring the former with techniques for evaluating the latter.

"The answer to anxieties over new technology is not to run 'tests' inconsistent with how the service is designed to be used, and to amplify the test's false and misleading conclusions through the news media," Matt Wood, general manager of artificial intelligence for Amazon's cloud-computing division, wrote in a January blog post. Amazon declined requests for an interview.

"I didn't know their reaction would be quite so hostile," Buolamwini said recently in an interview at her MIT lab.

Coming to her defense Wednesday was a coalition of researchers, including AI pioneer Yoshua Bengio, recent winner of the Turing Award, considered the tech field's version of the Nobel Prize.

They criticized Amazon's response, especially its distinction between facial recognition and analysis.



In this Feb. 22, 2019, photo, Washington County Sheriff's Office Deputy Jeff Talbot demonstrates how his agency used facial recognition software to help solve a crime at their headquarters in Hillsboro, Ore. The image on the left shows a man whose face was captured on a surveillance camera and investigators used the software to scan their database of past mug shots to match that facial image with an identity. (AP Photo/Gillian Flaccus)

"In contrast to Dr. Wood's claims, bias found in one system is cause for concern in the other, particularly in use cases that could severely impact people's lives, such as law enforcement applications," they wrote.

Its few publicly known clients have defended Amazon's system.

Chris Adzima, senior information systems analyst for the Washington County Sheriff's Office in Oregon, said the agency uses Amazon's

Recognition to identify the most likely matches among its collection of roughly 350,000 mug shots. But because a human makes the final decision, "the bias of that computer system is not transferred over into any results or any action taken," Adzima said.

But increasingly, regulators and legislators are having their doubts. A bipartisan bill in Congress seeks limits on facial recognition. Legislatures in Washington and Massachusetts are considering laws of their own.

Buolamwini said a major message of her research is that AI systems need to be carefully reviewed and consistently monitored if they're going to be used on the public. Not just to audit for accuracy, she said, but to ensure face recognition isn't abused to violate privacy or cause other harms.

"We can't just leave it to companies alone to do these kinds of checks," she said.

© 2019 The Associated Press. All rights reserved.

Citation: Face recognition researcher fights Amazon over biased AI (2019, April 3) retrieved 28 June 2024 from <https://phys.org/news/2019-04-recognition-amazon-biased-ai.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.