

TECHNOLOGY

Facial Recognition Is Accurate, if You're a White Guy

By STEVE LOHR FEB. 9, 2018

Facial recognition technology is improving by leaps and bounds. Some commercial software can now tell the gender of a person in a photograph.

When the person in the photo is a white man, the software is right 99 percent of the time.

But the darker the skin, the more errors arise — up to nearly 35 percent for images of darker skinned women, according to a new study that breaks fresh ground by measuring how the technology works on people of different races and gender.

These disparate results, calculated by Joy Buolamwini, a researcher at the M.I.T. Media Lab, show how some of the biases in the real world can seep into artificial intelligence, the computer systems that inform facial recognition.

In modern artificial intelligence, data rules. A.I. software is only as smart as the data used to train it. If there are many more white men than black women in the system, it will be worse at identifying the black women.

One widely used facial-recognition data set was estimated to be more than 75 percent male and more than 80 percent white, according to another research study.

The new study also raises broader questions of fairness and accountability in artificial intelligence at a time when investment in and adoption of the technology is racing ahead.

Today, facial recognition software is being deployed by companies in various ways, including to help target product pitches based on social media profile pictures. But companies are also experimenting with face identification and other A.I. technology as an ingredient in automated decisions with higher stakes like hiring and lending.

Researchers at the Georgetown Law School estimated that 117 million American adults are in face recognition networks used by law enforcement — and that African Americans were most likely to be singled out, because they were disproportionately represented in mug-shot databases.

Facial recognition technology is lightly regulated so far.

“This is the right time to be addressing how these A.I. systems work and where they fail — to make them socially accountable,” said Suresh Venkatasubramanian, a professor of computer science at the University of Utah.

Until now, there was anecdotal evidence of computer vision miscues, and occasionally in ways that suggested discrimination. In 2015, for example, Google had to apologize after its image-recognition photo app initially labeled African Americans as “gorillas.”

Sorelle Friedler, a computer scientist at Haverford College and a reviewing editor on Ms. Buolamwini’s research paper, said experts had long suspected that facial recognition software performed differently on different populations.

“But this is the first work I’m aware of that shows that empirically,” Ms. Friedler said.

Ms. Buolamwini, a young African-American computer scientist, experienced the bias of facial recognition firsthand. When she was an undergraduate at the Georgia Institute of Technology, programs would work well on her white friends, she said, but not recognize her face at all. She figured it was a flaw that would surely be fixed before long.

But a few years later, after joining the M.I.T. Media Lab, she ran into the missing-face problem again. Only when she put on a white mask did the software

recognize hers as a face.

By then, face recognition software was increasingly moving out of the lab and into the mainstream.

“O.K., this is serious,” she recalled deciding then. “Time to do something.”

So she turned her attention to fighting the bias built into digital technology. Now 28 and a doctoral student, after studying as a Rhodes scholar and a Fulbright fellow, she is an advocate in the new field of “algorithmic accountability,” which seeks to make automated decisions more transparent, explainable and fair.

Her short TED Talk on coded bias has been viewed more than 940,000 times, and she founded the Algorithmic Justice League, a project to raise awareness of the issue.

In her newly published paper, which will be presented at a conference this month, Ms. Buolamwini studied the performance of three leading face recognition systems — by Microsoft, IBM and Megvii of China — by classifying how well they could guess the gender of people with different skin tones. These companies were selected because they offered gender classification features in their facial analysis software — and their code was publicly available for testing.

She found them all wanting.

To test the commercial systems, Ms. Buolamwini built a data set of 1,270 faces, using faces of lawmakers from countries with a high percentage of women in office. The sources included three African nations with predominantly dark-skinned populations, and three Nordic countries with mainly light-skinned residents.

The African and Nordic faces were scored according to a six-point labeling system used by dermatologists to classify skin types. The medical classifications were determined to be more objective and precise than race.

Then, each company’s software was tested on the curated data, crafted for gender balance and a range of skin tones. The results varied somewhat. Microsoft’s error rate for darker-skinned women was 21 percent, while IBM’s and Megvii’s rates

were nearly 35 percent. They all had error rates below 1 percent for light-skinned males.

Ms. Buolamwini shared the research results with each of the companies. IBM said in a statement to her that the company had steadily improved its facial analysis software and was “deeply committed” to “unbiased” and “transparent” services. This month, the company said, it will roll out an improved service with a nearly 10-fold increase in accuracy on darker-skinned women.

Microsoft said that it had “already taken steps to improve the accuracy of our facial recognition technology” and that it was investing in research “to recognize, understand and remove bias.”

Ms. Buolamwini’s co-author on her paper is Timnit Gebru, who described her role as an adviser. Ms. Gebru is a scientist at Microsoft Research, working on its Fairness Accountability Transparency and Ethics in A.I. group.

Megvii, whose Face++ software is widely used for identification in online payment and ride-sharing services in China, did not reply to several requests for comment, Ms. Buolamwini said.

Ms. Buolamwini is releasing her data set for others to build upon. She describes her research as “a starting point, very much a first step” toward solutions.

Ms. Buolamwini is taking further steps in the technical community and beyond. She is working with the Institute of Electrical and Electronics Engineers, a large professional organization in computing, to set up a group to create standards for accountability and transparency in facial analysis software.

She meets regularly with other academics, public policy groups and philanthropies that are concerned about the impact of artificial intelligence. Darren Walker, president of the Ford Foundation, said that the new technology could be a “platform for opportunity,” but that it would not happen if it replicated and amplified bias and discrimination of the past.

“There is a battle going on for fairness, inclusion and justice in the digital world,” Mr. Walker said.

Part of the challenge, scientists say, is that there is so little diversity within the A.I. community.

“We’d have a lot more introspection and accountability in the field of A.I. if we had more people like Joy,” said Cathy O’Neil, a data scientist and author of “Weapons of Math Destruction.”

Technology, Ms. Buolamwini said, should be more attuned to the people who use it and the people it’s used on.

“You can’t have ethical A.I. that’s not inclusive,” she said. “And whoever is creating the technology is setting the standards.”

Correction: February 11, 2018

Because of an editing error, two picture captions with an earlier version of this article misstated the number of photos used by researchers to study artificial intelligence software’s accuracy in identifying two groups of people by gender. There were 385 photos of lighter-skinned men, not 296, and 296 of lighter-skinned women, not 385. Follow Steve Lohr on Twitter @SteveLohr

A version of this article appears in print on February 12, 2018, on Page B1 of the New York edition with the headline: Facial Recognition Works Best If You’re a White Guy.