


Study outlines what creates racial bias in facial recognition technology

December 7 2020, by Stephen Fontenot



Credit: CC0 Public Domain

As facial recognition technology comes into wider use worldwide, more attention has fallen on the imbalance in the technology's performance across races.

In a study published online Sept. 29 in the journal *IEEE Transactions on Biometrics, Behavior, and Identity Science*, researchers from The University of Texas at Dallas School of Behavioral and Brain Sciences (BBS) outlined the underlying factors that contribute to these deficits in facial recognition accuracy and offer a guide to assessing the algorithms as the technology improves.

Dr. Alice O'Toole, the Aage and Margareta Møller Professor in BBS, is the senior author of the study, which she describes as both "profound and unsatisfying" because it clarifies the scale of the challenge.

"Everybody's looking for a simple solution, but the fact that we outline these different ways that biases can happen—none of them being mutually exclusive—makes this a cautionary paper," she said. "If you're trying to fix an algorithm, be aware of how many different things are going on."

In a study last year conducted by the National Institute for Standards and Technology (NIST), the government agency found that the majority of facial recognition algorithms were far more likely to misidentify [racial minorities](#) than whites, with Asians, Blacks and Native Americans particularly at risk.

As a result of their research, the UT Dallas scientists concluded that while there isn't a one-size-fits-all solution for [racial bias](#) in facial recognition algorithms, there are specific approaches that can improve the technology's performance.

Psychological sciences doctoral student Jacqueline Cavazos, the study's lead author, divided the factors contributing to [bias](#) into two categories: data-driven and operationally defined. The former influence the algorithm's performance itself, while the latter originate with the user.

"Data-driven factors center on the most commonly theorized issues—that the training pool of images is in itself skewed," Cavazos said. "Are the images being used representative of groups? Are the training images of the same quality across races? Or is there something inherent about the algorithms' computation of face representations different between [race](#) groups?"

O'Toole added, "Our discussion of image difficulty for racial bias is a relatively new topic. We show that as pairs of images become more difficult to distinguish—as quality is reduced—racial bias becomes more pronounced. That hasn't been shown before."

Cavazos explained that operational bias can be introduced depending on where the threshold is set between matching and nonmatching decisions, and on what kinds of paired images are chosen.

"Our paper confirms what has been shown previously: Where you set the criterion for what is the same identity versus different identities can influence the error rate, and sometimes the same threshold will give you different error rates for different races," Cavazos said. "Secondly, you need to be sure that when you test an algorithm, pairs of images that are of different identities should always be matched on demographics;—this assures us that identification accuracy is based only on identity. Human participants are shown two images of different people with matching demographics—same race, same gender and so on. If algorithms aren't also presented in such similar pairs, algorithm performance can appear better than it really is because the machine's task is easier."

While the study outlines how racial bias should be evaluated in the use of facial recognition algorithms, the researchers emphasize that no easy solution to the problem exists.

"One of the novel things about this paper is how it brings all of these

factors together," O'Toole said. "Earlier work has centered on individual issues. But you have to look at them all to know the best way to use these algorithms."

O'Toole believes their research could help users understand which algorithms should be expected to show bias and how potentially to calibrate for that bias.

"For instance, you can measure the performance of an algorithm in a variety of ways. One measure might indicate that the algorithm is race biased, while another might not. Moreover, the algorithm could be biased in a way that you have not explicitly measured," O'Toole said.

"For example, one measure might be directly indicative of whether the algorithm could falsely accuse an innocent person. It might be aimed at determining how similar the people in two images have to appear for the machine to indicate that they are the same person. Another measure might focus on how many correct identifications the algorithm makes. These measures are made from the same algorithm, but they can easily dissociate, pointing to bias in one case and to equitable performance in the other case."

O'Toole said researchers in her field are still fighting myths that exist about facial recognition bias. One is the notion that bias is a problem unique to machines.

"Throwing the machines out the window won't make the process fair; humans have the same struggles that the algorithms do," she said.

Another challenge is overcoming the perception that race is an all-or-nothing descriptor.

"Race must not be viewed as categorical, or as if there's a finite list of races," O'Toole said. "In truth, biologically, race is continuous, so it's an

unreasonable expectation to think you can say 'race equity' and tune an algorithm for two races. This might disadvantage people of mixed race."

While the scientists agree that facial recognition algorithms have the potential to be helpful if one knows how to use them—and that newer algorithms have enabled significant progress against racial bias—they know there's a lot more work to be done.

"This is not to say that these algorithms shouldn't be used now as they currently are. But there are factors that need to be considered, and operation needs to be done with extreme caution," O'Toole said. "We have learned so much about the complexity of the problem that we have to acknowledge that there may never be a solution to the problem of making every face equally challenging to a face recognition [algorithm](#)."

More information: Jacqueline G. Cavazos et al. Accuracy comparison across face recognition algorithms: Where are we on measuring race bias?, *IEEE Transactions on Biometrics, Behavior, and Identity Science* (2020). [DOI: 10.1109/TBIOM.2020.3027269](https://doi.org/10.1109/TBIOM.2020.3027269)

Provided by University of Texas at Dallas

Citation: Study outlines what creates racial bias in facial recognition technology (2020, December 7) retrieved 20 March 2024 from <https://techxplore.com/news/2020-12-outlines-racial-bias-facial-recognition.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.
