



Observations

Can We Read a Person's Character from Facial Images?

The discredited “science” of physiognomy is making an unfortunate comeback

By Alexander Todorov on May 14, 2018




Credit: golubovy Getty Images

Physiognomy, the junk science of reading character from facial appearance, has a long history, with the first preserved documents referring to it dating back to Aristotle's time. Darwin almost missed his chance to take his historic voyage on the *Beagle* on account of his nose, because the captain of the ship—a fervent physiognomist—didn't believe that a person with such a nose would possess sufficient determination for the journey. "But I think," Darwin noted dryly in his autobiography, "he was afterwards well-satisfied that my nose had spoken falsely."

We may poke fun at the physiognomists' ideas, but the modern science of first impressions shows that we are all naïve physiognomists. We form instantaneous impressions of others from their facial appearance. Seeing a face for less than one tenth of a second is sufficient to make up our minds. First impressions are not only rapid but also consequential: We are more likely to vote for politicians who look competent; to invest in people who look trustworthy; and to assign harsher prison sentences to people who look the opposite. Faceism is a general feature of social life.

The modern science of first impressions has also identified many of the facial stereotypes that drive those impressions. In the last decade, psychologists have developed mathematical models that visualize these stereotypes. With these models, we can manipulate the appearance of faces by increasing or decreasing their perceived qualities of trustworthiness and competence as we desire. And more importantly, we can build and test theories about the origins of facial stereotypes.



ADVERTISEMENT


However, one of the unintended consequences of the progress in this research has been the revival of physiognomy. Perhaps our facial stereotypes are not just stereotypes, but a true window into the character of others. Correspondingly, there has been a proliferation of studies claiming that we can discern all kinds of private things about others such as their mental health, political and sexual orientation, and so on, from their facial images alone.

These claims are typically based on the finding that human guesses about, say, sexual orientation, are better than chance. The problem is that these guesses are barely better than chance and often less accurate than guesses based on more general knowledge.

Moreover, many of these studies are based on the fallacy that all facial images are equally representative of the face's owner. While this assumption may ring true in the case of familiar faces, which are easily recognizable from different images, it is certainly false in the case of unfamiliar faces—and by definition, first impressions are about unfamiliar faces. Often, we cannot tell whether two different images represent the same (unfamiliar) person, and these images can trigger completely different impressions. Hence, how the images are sampled is a critical issue when assessing the accuracy of first impressions.

Consider how biases in the sampling of images can affect inferences about the accuracy of first impressions. In many “gaydar” studies, participants guess the sexual orientation of others from images posted on online dating websites. In one of the very first such studies, the guesses were accurate about 58 percent of the time (where chance is 50 percent). But since we strategically select the images we post to represent ourselves to the kinds of people we want to attract, this isn’t a neutral sample.

In fact, when the guesses were based on online images of gay and heterosexual men posted by their friends (far from a perfect control), they were only accurate 52 percent of the time. This kind of result isn’t only true when subjects are guessing sexual orientation. In a recent study, the researchers used images from online dating websites to test whether participants can guess social class, represented by wealth. The participants were accurate about 57 percent of the time. But when the guesses were based on images taken under standardized conditions, the accuracy dropped to 51.5 percent.



ADVERTISEMENT

With the ubiquity of online face images, studies attempting to read our “essence” from these images are not going to disappear. In the last few years, there has been a new wave of artificial intelligence


(AI) studies attempting to do exactly that. A technology start-up is already offering facial profiling services to private companies and governments. Last year, two computer scientists posted a non-peer-reviewed paper online claiming that their algorithm can guess the criminality of people from a single facial image. And in the last month, a prestigious journal accepted for publication a paper claiming that AI algorithms can detect sexual orientation from facial images with seemingly surprising accuracy.

However, the same problems that apply to human studies apply to AI studies as well. The latter use powerful algorithms that can detect subtle but systematic differences between two sets of images. But the sample of images used to train the algorithm is just as important as the algorithm itself. In the paper on criminality, the authors provided a few images of “criminals” and “non-criminals.” Besides obvious differences in facial expressions, the “criminals” wore t-shirts while the “non-criminals” wore suits. A powerful algorithm would easily pick up these differences and produce a seemingly accurate classification.

The fallacy that all facial images are equally representative of the face owner plays in even more subtle ways in AI studies, especially when the claim is that the algorithms are measuring invariant facial features from 2-D images. Camera-to-head distance, camera parameters, slight head tilts, subtle expressions and many other apparently trivial differences affect the measurement of what are meant to be stable morphological features. When these differences are not controlled for, the AI studies simply amplify our human biases.


Moreover, the implications of using AI to do “face reading” are morally abhorrent. The senior author of the sexual orientation paper claims that his main motivation was to warn the LGBT community about the potential of this technology to harm them, especially in repressive countries. But while the study claims to identify real morphological differences between gay and straight people, all it really shows is that an algorithm can identify openly gay people from their self-posted images—just as ordinary humans can.

This is precisely the kind of “scientific” claim that can motivate repressive governments to apply AI algorithms to images of their citizens. And what is it to stop them from “reading” intelligence, political orientation and criminal inclinations from these images?



ADVERTISEMENT

The views expressed are those of the author(s) and are not necessarily those of Scientific American.



ADVERTISEMENT

ABOUT THE AUTHOR(S)

Alexander Todorov

Alexander Todorov is a professor of psychology at Princeton University and author of *Face Value: The Irresistible Influence of First Impressions* (Princeton University Press, 2017).

Scientific American is part of Springer Nature, which owns or has commercial relations with thousands of scientific publications (many of them can be found at www.springernature.com/us). Scientific American maintains a strict policy of editorial independence in reporting developments in science to our readers.

© 2018 SCIENTIFIC AMERICAN, A DIVISION OF NATURE AMERICA, INC.

ALL RIGHTS RESERVED.