

Supposedly gay-identifying AI tells us more about stereotypes than the origins of sexuality

September 14 2017, by Colin Klein



Credit: CC0 Public Domain

In a [forthcoming paper](#), two Stanford researchers used a deep neural network to detect sexuality from profile pictures on a US dating website.

The internet [was aghast](#). The authors themselves [raised the spectre](#) of Orwellian surveillance.

More problematic, however, was their claim that the results provide support for a controversial theory that broadly suggests gay people appear and act atypical for their gender.

This conclusion threatens to undermine science with stereotype.

How did the study work?

Deep neural networks are extremely powerful tools. They are especially good at tasks like classifying pictures, as they can combine innumerable subtle cues that are difficult for humans to register.

The study itself avoids obvious pitfalls in its method, though it was explicitly limited in scope given the researchers used only clear pictures of Caucasians.

Their classifier was given two faces, one gay and one straight, and asked to say which was which. This is a much easier task than classifying a single face (imagine if baggage handlers only ever had to pick which of two x-rays definitely had a gun).

Much of the immediate criticism focused on these limitations. Yet easier tasks are often enough for proof of concept, and the network achieved a good level of accuracy.

The "prenatal hormone theory" of sexuality

Interpreting that accuracy is where things get dicey.

The authors suggest the results affirm the "prenatal [hormone](#) theory" of sexuality, which claims that atypical hormone exposure in the womb gives gay men more "feminine" brains (similarly for lesbians and masculinity).

As they put it:

"Our results provide strong support for the [Prenatal Hormone Theory], which argues that same-gender sexual orientation stems from the underexposure of male fetuses and overexposure of female fetuses to prenatal androgens responsible for the sexual differentiation of faces, preferences, and behavior."

Some [older work](#) on the hypothesis was meant to be politically progressive. Yet the prenatal hormone theory has attracted criticism, not least because it seems to revive outdated stereotypes of "prancing queens" and "butch dykes".

The best evidence for prenatal influence on gendered behaviour depends [on studies](#) of children with abnormalities in hormone sensitivity. Even in those extreme cases, the evidence suggests [an inconsistent](#) link between prenatal hormones and either gendered behaviour or sexual identity.

Indeed, the evidence for robust "masculine" and "feminine" brain differences is not clear. As psychologist [Cordelia Fine](#) details in her recent book *Testosterone Rex*, the differences between male and female brains are small compared to the enormous variability within genders.

The brain also [changes dramatically](#) in response to the environment. Neuroscientist [Lise Eliot](#) points out that very small neural differences in gender can be magnified by the different social worlds in which men and women live.

What did the study actually show?

Should this particular neural network change our minds?

The authors' logic is that prenatal hormone differences would affect both facial structure and sexual preference. Thus a woman with a more "male-typical" facial structure is also more likely to have a "male-typical" preference for female partners.

Since, according to this paper, sexuality can be predicted from facial information, that must give some evidence in favour of the prenatal hormone theory.

Yet as my colleagues and I have [recently argued](#), the results spit out by neural networks can be notoriously difficult to interpret. Their power makes it hard to know how high accuracy is achieved.

The authors tried to sort out which features were especially important to the network. Some are about basic face shape. Others aren't. As the authors note in their [response to critics](#):

"The gender atypicality of gay faces extended beyond morphology. Lesbians tended to use less eye makeup, had darker hair, and wore less revealing clothes... "

Similarly, gay men are less likely to have facial hair, "lesbians tended to wear baseball caps", and numerous other differences. These are not features that are straightforwardly determined by prenatal hormones. (Astute readers will note that some aren't properties of faces at all.)

Further, these are only the differences that were obvious to the

researchers. Deep neural networks commonly rely on feature combinations that are meaningless to humans.

Stereotypes versus science

But now we get to a place where politically powerful stereotypes undermine science.

The argument supposes that lesbians are less "feminine" because they don't match the grooming style of straight women. This only works if straight women are defined as gender typical.

This is a case of what queer theorists term [heteronormativity](#): the idea that heterosexuality is the norm, and differences from it must be deviant and require special explanation.

Yet analogous arguments are obviously preposterous. Suppose the researchers found that straight [black men](#) tend to have shorter hair than white straight men. Would we conclude that black men are gender-atypical? Absurd. Black male hairstyles are gender typical for black men.

Similarly, lesbian grooming styles are gender typical for lesbian women. Equally so for all of the innumerable subtle signs and signals that we learn to use to convey all aspects of our lives.

The rise of LGBTQ visibility in the past two decades has done much to combat both heteronormative assumptions and the outdated stereotypes upon which the prenatal hormone theory rests.

The impact of dubious prejudices

Overlooking this variation is a classic example of how [stereotypes can distort thinking](#).

Of course, the authors of the study mostly dismiss this. They claim that "We also know many very old men, which does not invalidate the statement that women tend to live longer." Yet the gender differences in death rates are [large and visible to the naked eye](#).

The whole point of using a neural network, recall, was to pick up on numerous tiny differences which humans cannot register. One possible conclusion to draw is that any underlying biological differences are similarly minuscule.

If the differences are that small, their scientific relevance is doubtful: either prenatal hormones don't differ that much, or the effect that they do have is tiny.

Much of the subsequent discussion around the paper focuses on the ethics of automatically detecting sexual preference. As the writer Jaron Lanier notes, most apocalyptic scenarios involving AI tend to obscure the very real role AI can play in [perpetuating entrenched disadvantage](#).

In a world without prejudice against the LGBTQ community, the ability to detect sexuality from a photograph would be ethically neutral. We have far more to worry about from outdated science that embodies dubious prejudices than we do from deep learning networks.

This article was originally published on [The Conversation](#). Read the [original article](#).

Provided by The Conversation

Citation: Supposedly gay-identifying AI tells us more about stereotypes than the origins of sexuality (2017, September 14) retrieved 6 May 2024 from <https://phys.org/news/2017-09-supposedly-gay-identifying-ai-stereotypes-sexuality.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.