

Can you spot the AI impostors? Research finds AI faces can look more real than actual humans

November 14 2023, by Amy Dawel et al



Credit: Pixabay/CC0 Public Domain

Does ChatGPT ever give you the eerie sense you're interacting with another human being?



Artificial intelligence (AI) has reached an astounding level of realism, to the point that some tools can even <u>fool people</u> into thinking they are interacting with <u>another human</u>.

The eeriness doesn't stop there. In a <u>study published today</u> in *Psychological Science*, we've discovered images of white faces generated by the popular <u>StyleGAN2 algorithm</u> look more "human" than actual people's faces.

AI creates hyperrealistic faces

For our research, we showed 124 participants pictures of many different white faces and asked them to decide whether each face was real or generated by AI.

Half the pictures were of real faces, while half were AI-generated. If the participants had guessed randomly, we would expect them to be correct about half the time—akin to flipping a coin and getting tails half the time.

Instead, participants were systematically wrong, and were more likely to say AI-generated faces were real. On average, people labeled about 2 out of 3 of the AI-generated faces as human.

These results suggest AI-generated faces look more real than actual faces; we call this effect "hyperrealism." They also suggest people, on average, aren't very good at detecting AI-generated faces. You can compare for yourself the portraits of real people at the top of the page with the ones embedded below.

But perhaps people are aware of their own limitations, and therefore aren't likely to fall prey to AI-generated faces online?



To find out, we asked participants how confident they felt about their decisions. Paradoxically, the people who were the worst at identifying AI impostors were the most confident in their guesses.

In other words, the people who were most susceptible to being tricked by AI weren't even aware they were being deceived.

Biased training data deliver biased outputs

The <u>fourth industrial revolution</u>—which includes technologies such as AI, robotics and advanced computing—has profoundly changed the kinds of "faces" we see online.

AI-generated faces are readily available, and their use comes with both risks and benefits. Although they have been used to <u>help find missing</u> <u>children</u>, they have also been used in <u>identity fraud</u>, <u>catfishing</u> and <u>cyber</u> <u>warfare</u>.

People's misplaced confidence in their ability to detect AI faces could make them more susceptible to deceptive practices. They may, for instance, readily hand over sensitive information to cybercriminals masquerading behind hyperrealistic AI identities.

Another worrying aspect of AI hyperrealism is that it's racially biased. Using <u>data from another study</u> which also tested Asian and Black faces, we found only white AI-generated faces looked hyperreal.

When asked to decide whether faces of color were human or AIgenerated, participants guessed correctly about half the time—akin to guessing randomly.

This means white AI-generated faces look more real than AI-generated faces of color, as well as white human faces.



Implications of bias and hyperrealistic AI

This <u>racial bias</u> likely stems from the fact that AI algorithms, including the one we tested, are often trained on images of mostly white faces.

Racial bias in algorithmic training can have serious implications. One recent study found self-driving cars are <u>less likely to detect Black people</u>, placing them at greater risk than white people. Both the companies producing AI, and the governments overseeing them, have a responsibility to ensure diverse representation and mitigate bias in AI.

The realism of AI-generated content also raises questions about our ability to accurately detect it and protect ourselves.

In our research, we identified several features that make white AI faces look hyperreal. For instance, they often have proportionate and familiar features, and they lack distinctive characteristics that make them stand out as "odd" from other faces. Participants misinterpreted these features as signs of "humanness," leading to the hyperrealism effect.

At the same time, AI technology is advancing so rapidly it will be interesting to see how long these findings apply. There's also no guarantee AI faces generated by other algorithms will differ from human faces in the same ways as those we tested.

Since our study was published, we have also tested the ability of AI detection technology to identify our AI faces. Although this technology claims to identify the particular type of AI faces we used with a high accuracy, it performed as poorly as our human participants.

Similarly, software for detecting AI writing has also had <u>high rates</u> of <u>falsely accusing people of cheating</u>—especially people whose native language is not English.



Managing the risks of AI

So how can people protect themselves from misidentifying AI-generated content as real?

One way is to simply be aware of how poorly people perform when tasked with separating AI-generated faces from real ones. If we are more wary of our own limitations on this front, we may be less easily influenced by what we see online—and can take additional steps to verify information when it matters.

Public policy also plays an important role. One option is to require the use of AI to be declared. However, this might not help, or may inadvertently provide a false sense of security when AI is used for deceptive purposes—in which case it is almost impossible to police.

Another approach is to focus on authenticating trusted sources. Similar to the "Made in Australia" or "European CE tag," applying a trusted source badge—which can be verified and has to be earned through rigorous checks—could help users select reliable media.

This article is republished from <u>The Conversation</u> under a Creative Commons license. Read the <u>original article</u>.

Provided by The Conversation

Citation: Can you spot the AI impostors? Research finds AI faces can look more real than actual humans (2023, November 14) retrieved 21 May 2024 from https://medicalxpress.com/news/2023-11-ai-impostors-real-actual-humans.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.