How to spot deepfakes? Look at light reflection in the eyes

March 11 2021, by Melvin Bankhead Iii

Question: Which of these people are fake? Answer: All of them. Credit: thispersondoesnotexist.com and the University at Buffalo.

University at Buffalo computer scientists have developed a tool that automatically identifies deepfake photos by analyzing light reflections in the eyes.

The tool proved 94% effective with portrait-like photos in experiments described in a paper accepted at the IEEE International Conference on Acoustics, Speech and Signal Processing to be held in June in Toronto, Canada.

"The cornea is almost like a perfect semisphere and is very reflective," says the paper's lead author, Siwei Lyu, Ph.D., SUNY Empire Innovation Professor in the Department of Computer Science and Engineering. "So, anything that is coming to the eye with a light emitting from those sources will have an image on the cornea.

"The two eyes should have very similar reflective patterns because they're seeing the same thing. It's something that we typically don't typically notice when we look at a face," says Lyu, a multimedia and digital forensics expert who has testified before Congress.

The paper, "Exposing GAN-Generated Faces Using Inconsistent Corneal Specular Highlights," is available on the open access repository arXiv.

Co-authors are Shu Hu, a third-year computer science Ph.D. student and research assistant in the Media Forensic Lab at UB, and Yuezun Li, Ph.D., a former senior research scientist at UB who is now a lecturer at the Ocean University of China's Center on Artificial Intelligence.

Tool maps face, examines tiny differences in eyes

When we look at something, the image of what we see is reflected in our eyes. In a real photo or video, the reflections on the eyes would generally appear to be the same shape and color.

However, most images generated by artificial intelligence—including generative adversary network (GAN) images—fail to accurately or consistently do this, possibly due to many photos combined to generate

the fake image.

Lyu's tool exploits this shortcoming by spotting tiny deviations in reflected light in the eyes of deepfake images.

To conduct the experiments, the research team obtained real images from Flickr Faces-HQ, as well as fake images from www.thispersondoesnotexist.com, a repository of AI-generated faces that look lifelike but are indeed fake. All images were portrait-like (real people and fake people looking directly into the camera with good lighting) and 1,024 by 1,024 pixels.

The tool works by mapping out each face. It then examines the eyes, followed by the eyeballs and lastly the light reflected in each eyeball. It compares in incredible detail potential differences in shape, light intensity and other features of the reflected light.

"Deepfake-o-meter," and commitment to fight deepfakes

While promising, Lyu's technique has limitations.

For one, you need a reflected source of light. Also, mismatched light reflections of the eyes can be fixed during editing of the image. Additionally, the technique looks only at the individual pixels reflected in the eyes—not the shape of the eye, the shapes within the eyes, or the nature of what's reflected in the eyes.

Finally, the technique compares the reflections within both eyes. If the subject is missing an eye, or the eye is not visible, the technique fails.

Lyu, who has researched machine learning and computer vision projects

for over 20 years, previously proved that deepfake videos tend to have [inconsistent or nonexistent blink rates](#) for the video subjects.

In addition to testifying before Congress, he [assisted Facebook](#) in 2020 with its deepfake detection global challenge, and he helped create the "[Deepfake-o-meter](#)," an online resource to help the average person test to see if the video they've watched is, in fact, a deepfake.

He says identifying deepfakes is increasingly important, especially given the hyper-partisan world full of race-and gender-related tensions and the dangers of disinformation—particularly violence.

"Unfortunately, a big chunk of these kinds of fake videos were created for pornographic purposes, and that (caused) a lot of … psychological damage to the victims," Lyu says. "There's also the potential political impact, the fake video showing politicians saying something or doing something that they're not supposed to do. That's bad."

More information: Exposing GAN-generated Faces Using Inconsistent Corneal Specular Highlights. arXiv:2009.11924v2 [cs.CV] arxiv.org/abs/2009.11924

Provided by University at Buffalo