Nvidia's face-making approach is genuinely GAN-tastic

December 19 2018, by Nancy Cohen

Uncurated set of images produced by the style-based generator (config F) with the FFHQ dataset. Credit: arXiv:1812.04948 [cs.NE]

A new kind of Generative Adversarial Network approach has technology observers scratching their heads: How can images be fake and yet look so real?

"We came up with a new generator that automatically learns to separate different aspects of the images without any human supervision," the researchers said in a video. They stated in their paper, "The new architecture leads to an automatically learned, unsupervised separation of high-level attributes."

Making fake pictures look real is an artistic effort that is hardly new, but these three researchers raised the effort to the next level.

They spelled out their technique in their paper, "A Style-Based Generator Architecture for Generative Adversarial Networks." The paper is on the arXiv and it has attracted much attention.

Stephen Johnson in *Big Think* said the results were "pretty startling." Will Knight in *MIT Technology Review* said what we are looking at constitutes "stunning, almost eerie realism."

The researchers, Tero Karras, Samuli Laine, and Timo Aila, are from Nvidia. Their approach focuses on constructing a generative adversarial network, or GAN, where learning takes place to generate entirely new images that *mimic* the appearance of real photos.

The authors said all images in this video were produced by their generator. "They are not photographs of real people."

Their paper," said Knight, showed how the approach can be used to play with, and remix elements such as race, gender, or even freckles.

The magic sauce is their style-based generator. *Big Think* explained this

as a modified version of the conventional technology that is used to automatically generate images.

Their technology is messing with your head and having the last laugh (or word, however you look at it).

You as a human think "images." The generator, however, thinks "style collection."

Visualizing the effect of styles in the generator by having the styles produced by one latent code (source) override a subset of the styles of another one (destination). Credit: arXiv:1812.04948 [cs.NE]

Each style controls effects at a particular scale. There are coarse styles,

middle styles, fine styles. (Coarse styles refer to pose, hair, face shape; middle styles refer to facial features; eyes. Fine styles refer to the color scheme.)

Will Knight, meanwhile, made some observations about GANs: "GANs employ two dueling neural networks to train a computer to learn the nature of a data set well enough to generate convincing fakes. When applied to images, this provides a way to generate often highly realistic fakery."

Johnson provided a background of the GAN concept:

"In 2014, a researcher named Ian Goodfellow and his colleagues wrote a paper outlining a new machine learning concept called generative adversarial networks. The idea, in simplified terms, involves pitting two neural networks against each other. One acts as a generator that looks at, say, pictures of dogs and then does its best to create an image of what it thinks a dog looks like. The other network acts as a discriminator that tries to tell fake images from real ones.

"At first, the generator might produce some images that don't look like dogs, so the discriminator shoots them down. But the generator now knows a bit about where it went wrong, so the next image it creates is slightly better. This process continues until, in theory, the generator creates a good image of a dog."

Nvidia's team added style transfer principles to the GAN mix.

Devin Coldewey in *TechCrunch*: "Cars, cats, landscapes—all this stuff more or less fits the same paradigm of small, medium and large features that can be isolated and reproduced individually."

From a technical view, their work has been praised for impressive results

in images of people who look real. From a popular view, sharpened by talk of fake news, the advancement is seen as dangerous. "The ability to generate realistic artificial images, often called deepfakes when images are meant to look like recognizable people, has raised concern in recent years," said Johnson.

Tech-watching site pages were filled with comments about how "creepy" this was, and some comments were simply questions: "Why are we doing this?" "Who is paying for it?" "If we don't create regulatory boundaries, I think the next major battle of human history will be fought (and possibly lost) against the AI we create," said one response.

Nonetheless, not all comments reflected unease. This is a technology advancement and some comments pointed out that applications could be helpful in certain sectors. Designers, advertising agency creators, and even video game creators could use this technology step forward.

"These AI-generated faces promise to usher in a new generation of photorealistic people in video games and movies without the need for human actors or extras," said *HotHardware*.

  More information: A Style-Based Generator Architecture for Generative Adversarial Networks arXiv:1812.04948 [cs.NE] arxiv.org/abs/1812.04948