


Artificial intelligence makes blurry faces look more than 60 times sharper

June 12 2020, by Robin A. Smith



This AI turns blurry pixelated photos into hyperrealistic portraits that look like real people. The system automatically increases any image's resolution up to 64x, 'imagining' features such as pores and eyelashes that weren't there in the first place. Credit: Duke University

Duke University researchers have developed an AI tool that can turn

blurry, unrecognizable pictures of people's faces into eerily convincing computer-generated portraits, in finer detail than ever before.

Previous methods can scale an image of a face up to eight times its original resolution. But the Duke team has come up with a way to take a handful of pixels and create realistic-looking faces with up to 64 times the resolution, 'imagining' features such as fine lines, eyelashes and stubble that weren't there in the first place.

"Never have super-resolution images been created at this resolution before with this much detail," said Duke computer scientist Cynthia Rudin, who led the team.

The system cannot be used to identify people, the researchers say: It won't turn an out-of-focus, unrecognizable photo from a security camera into a crystal clear image of a real person. Rather, it is capable of generating new faces that don't exist, but look plausibly real.

While the researchers focused on faces as a proof of concept, the same technique could in theory take low-res shots of almost anything and create sharp, realistic-looking pictures, with applications ranging from medicine and microscopy to astronomy and satellite imagery, said co-author Sachit Menon '20, who just graduated from Duke with a double-major in mathematics and computer science.



Facial features like eyes and lips are barely distinguishable in the blurry photo on the left. Enlarged more than 60 times (right) it's a different story—thanks to artificial intelligence. Credit: Duke University

The researchers will present their method, called PULSE, at the 2020 Conference on Computer Vision and Pattern Recognition (CVPR), held virtually from June 14 to June 19.


Traditional approaches take a low-resolution image and 'guess' what extra pixels are needed by trying to get them to match, on average, with corresponding pixels in high-resolution images the computer has seen before. As a result of this averaging, textured areas in hair and skin that might not line up perfectly from one pixel to the next end up looking fuzzy and indistinct.

The Duke team came up with a different approach. Instead of taking a low-resolution image and slowly adding new detail, the system scours AI-

generated examples of high-resolution faces, searching for ones that look as much as possible like the input image when shrunk down to the same size.

The team used a tool in machine learning called a "[generative adversarial network](#)," or GAN, which are two [neural networks](#) trained on the same data set of photos. One network comes up with AI-created human faces that mimic the ones it was trained on, while the other takes this output and decides if it is convincing enough to be mistaken for the real thing. The first network gets better and better with experience, until the second network can't tell the difference.

PULSE can create realistic-looking images from noisy, poor-quality input that other methods can't, Rudin said. From a single blurred image of a face it can spit out any number of uncannily lifelike possibilities, each of which looks subtly like a different person.



Meet the authors: Sachit Menon, Alex Damian, McCourt Hu, Nikhil Ravi and Cynthia Rudin. From a single blurred image PULSE can generate uncannily

lifelike portraits, which might differ subtly from the real person but are much sharper than previous methods. Credit: Duke University

Even given pixelated photos where the eyes and mouth are barely recognizable, "our algorithm still manages to do something with it, which is something that traditional approaches can't do," said co-author Alex Damian '20, a Duke math major.

The system can convert a 16x16-[pixel](#) image of a face to 1024 x 1024 pixels in a few seconds, adding more than a million pixels, akin to HD resolution. Details such as pores, wrinkles, and wisps of hair that are imperceptible in the low-res photos become crisp and clear in the computer-generated versions.

The researchers asked 40 people to rate 1,440 images generated via PULSE and five other scaling methods on a scale of one to five, and PULSE did the best, scoring almost as high as high-quality photos of actual people.

See the results and upload images for yourself at <http://pulse.cs.duke.edu/>.

More information: PULSE: Self-Supervised Photo Upsampling via Latent Space Exploration of Generative Models, arXiv:2003.03808 [cs.CV] arxiv.org/abs/2003.03808

Provided by Duke University

Citation: Artificial intelligence makes blurry faces look more than 60 times sharper (2020, June

12) retrieved 26 April 2024 from <https://techxplore.com/news/2020-06-artificial-intelligence-blurry-sharper.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.