REPORT TECH ARTIFICIAL INTELLIGENCE

Artificial intelligence is going to make it easier than ever to fake images and video

by James Vincent | Oec 20, 2016, 10:43am EST

SHARE TWEET LINKEDIN

Asample of Smile Vector's work Tom White

Smile Vector is a Twitter bot that can <u>make any celebrity smile</u>. It scrapes the web for pictures of faces, and then it morphs their expressions using a deep-learning-powered neural network. Its results aren't perfect, but they're created completely automatically, and it's just a small hint of what's to come as artificial intelligence opens a new world of image, audio, and video fakery. Imagine a version of Photoshop that can edit an image as easily as you can edit a Word document — will we ever trust our own eyes again?

"THIS WILL BE A QUANTUM STEP FORWARD."

"I definitely think that this will be a quantum step forward," Tom White, the creator of Smile Vector, tells *The Verge*. "Not only in our ability to manipulate images but really their prevalence in our society." White says he created his bot in order to be "provocative," and to show people what's happening with AI in this space. "I don't think many people outside the machine learning community knew this was even possible," says White, a lecturer in creative coding at Victoria University School of design. "You can imagine an Instagram-like filter that just says 'more smile' or 'less smile,' and suddenly that's in everyone's pocket and everyone can use it."

Smile Vector is just the tip of the iceberg. It's hard to give a comprehensive overview of all the work being done on multimedia manipulation in AI right now, but here are a few examples: creating 3D face models from a single 2D image; changing the facial expressions of a target on video in realtime using a human "puppet"; changing the light source and shadows in any picture; generating sound effects based on mute video; live-streaming the presidential debates but making Trump bald; "resurrecting" Joey from Friends using old clips; and so on. Individually, each of these examples is a curiosity; collectively, they add up to a whole lot more.

"The field is progressing extremely rapidly," says Jeff Clune, an assistant professor of computer science at the University of Wyoming. "Jaw-dropping examples arrive in my inbox every month." Clune's own work isn't about manipulating images, but generating them, whole cloth. His team at Wyoming began work on this in 2015 by adapting neural networks trained in object recognition. Inspired by research done on the human brain in 2005, they identified the neurons that lit up when faced with certain images, and taught the network to produce the images that maximized this stimulation.

In 2015, their networks were creating pictures like this:





Jason Yosinski, Jeff Clune, Anh Nguyen, Thomas Fuchs, and Hod Lipson

In 2016, they're creating pictures like this:





Anh Nguyen, Jason Yosinski, Yoshua Bengio, Alexey Dosovitskiy, Jeff Clune

To create these images, the neural network is trained on a database of similar pictures. Then, once it's absorbed enough images of ants, redshanks, and volcanoes it can produce its own versions on command — no instruction other than "show me a volcano" is needed. The two bottlenecks at the moment are image resolution (none of these pictures are bigger than 256 x 256) and finding the labeled pictures to train the networks with. "Our current limitation isn't the capability of the models but the existence of data sets at higher resolution," says Clune. "How long until we can produce full HD images that are photorealistic? It's anyone's guess, but it's probably on the order of years, not decades."

STYLE TRANSFER WENT FROM RESEARCH PAPER TO TOP APP IN LESS THAN A YEAR

Once these techniques have been perfected, they spread quickly. A good example is a method known as "style



transfer" which uses neural networks to apply the characteristics of one image to another. A key paper on this subject was published in September 2015, with researchers turning this work into an open-source web app in January 2016. In June, a Russian startup finessed this code into a mobile app named Prisma, which allowed anyone to apply various art styles to pictures on their phones and share them in various social networks. The app exploded in popularity, and this November, Facebook unveiled its own version, adding a couple of new features along the way. From cutting-edge research to commercial product in less than a year; that's how fast these tools can be adopted.

Clune says that in the future, Al-powered image generation will be useful in the creative industries. A furniture designer could use it as an "intuition pump," he says, feeding a generative network a database of chairs, and then asking it to generate its own variants which the designer could perfect. Another use might be creating content for video games and virtual reality, which users could literally dictate in real time. Want dragons? Just ask for them. Want bigger dragons, with guns for arms and bright purple mohawks? Also fine. Researchers are already working on precursors to this type of interface. In the picture below, the images on the right were created based on the captions on the left, nothing more.





Scott Reed, Zeynep Akata, Xinchen Yan, Lajanugen Logeswaran, Bernt Schiele, Honglak Lee

Another obvious beneficiary would be hoaxes. Consider the video below — a demonstration of a program called Face2Face, which essentially turns people into puppets, letting you map their facial expression to your own. The researchers demonstrate it using footage of Trump and Obama. Now combine that with prototype software recently unveiled by Adobe that lets you edit human speech (the company says it could be used for fixing voiceovers and dialog in films). Then you can create video footage of politicians, celebrities, saying, well, whatever you want them, too. Post your clip on any moderately popular Facebook page, and watch it spread around the internet.



That's not to say these tools will steer society into some fact-less free-for-all. After all, the practice of retouching photos goes all the way back to the dark room, and the media has often been tricked into reporting fakes images as real. Anything from North Korean "missile launches" to pictures of Osama bin Laden's "corpse" splashed on the pages of British tabloids. And the same can be done with video — see, for example, the 2015 Planned Parenthood scandal that relied on undercover footage that had been edited to support sensational and false claims.

However, we can't deny that digital tools will allow more people to create these sorts of fakes. There's nothing Al can do to an image or a video that a human expert couldn't (given enough time), but once *everyone* can doctor a photo as easily as creating a Word document, it would be overly optimistic to claim there won't be any side effects. Al-powered fakes and manipulations aren't hard to spot now (blurring is one of the most common tells, as is low resolution and just plain "looking fake"), but researchers say they're just going to get better and better.

WHAT HAPPENS WHEN EVERYONE CAN DOCTOR A PHOTO AS QUICKLY AND EASILY AS A PROFESSIONAL?

The proliferation of realistic fakes would be a boon to conspiracy theorists, and would contribute to the current climate of deteriorating confidence in journalism. Once people know there are fake images being circulated, it gives them reason to doubt real images they might not want to believe, for whatever reason. (See, for example, this 2012 blog of Hurricane Sandy photos, which not only verifies fakes, but also genuine images.) And if new software allows us to manipulate audio and video content as easily as images, it would undermine another pillar of "reliable" evidence.

Al researchers involved in this fields are already getting a firsthand experience of the coming media environment. "I currently exist in a world of reality vertigo," says Clune. "People send me real images and I start to wonder if they look fake. And when they send me fake images I assume they're real *because the quality is so good*. Increasingly, I think, we won't know the difference between the real and the fake. It's up to people to try and educate themselves."

MORE FIRHOEM VERGE

Apple AirPods review: wireless that wows, earbuds that don't

Web2PDF

Tim Cook says 'great desktop' Macs are in the works

Rogue One director says its original ending was very different

A filmmaker installed security software on a decoy phone to spy on smartphone thieves

Question Club: We throw down over Rogue One's CGI characters, choppy first act, and that damn blue milk

RECOMMENDED

Recommended by **outbrain**



Dell goes wide with new UltraSharp 34 professional monitor

Sponsored | Digital Trends



Do This Every Day And You'll Never Need a New Computer

Sponsored | Smart Web User

Get an interactive tour to wheat breeding worldwide

Sponsored | Bayer Crop Science





Lenovo made a very 2016 smartphone, just in time for 2017



FCC Chairman Tom Wheeler to step down after Trump's inauguration



Trump fires transition team member for spreading Pizzagate conspiracy theory

Loading comments...

T H E L A T E S T

Literally sticking Apple AirPods in your ear is one way to stop them from falling out

by Rich McCormick

You can now buy refurbished Apple Watch models for as little as \$229

by Nick Statt | @nickstatt

0

President Obama blocks new offshore drilling in Arctic and Atlantic

by Rachel Becker





A Google employee is suing the company for being too confidential

by Nick Statt | @nickstatt



Congressional group says backdoor laws would do more harm than good

by Russell Brandom | @russellbrandom





IMAX's VR theaters aren't coming this year after all

by Adi Robertson | @thedextriarchy



All Systems Operational .Check out our status page for more details.

VOX MEDIA

Advertise with us
Jobs @ Vox Media
© 2016 Vox Media, Inc. All Rights Reserved

