

Teaching artificial intelligence to create visuals with more common sense

July 1 2019, by Adam Conner-Simons



The GANpaint system developed at MIT can easily add features to an existing image. At left, the original photo of a kitchen; at right, the same kitchen with the addition of a window. Co-author Jun-Yan Zhu believes better understanding of GANs will help researchers be able to better stamp out fakery: “This understanding may potentially help us detect fake images more easily.” Credit: Massachusetts Institute of Technology

Today's smartphones often use artificial intelligence (AI) to help make the photos we take crisper and clearer. But what if these AI tools could be used to create entire scenes from scratch?

A team from MIT and IBM has now done exactly that with "GANpaint Studio," a system that can automatically generate realistic photographic images and edit objects inside them. In addition to helping artists and designers make quick adjustments to visuals, the researchers say the work may help computer scientists identify "fake" images.

David Bau, a Ph.D. student at MIT's Computer Science and Artificial Intelligence Lab (CSAIL), describes the project as one of the first times [computer scientists](#) have been able to actually "paint with the neurons" of a neural network—specifically, a popular type of network called a generative adversarial network (GAN).

Available online as [an interactive demo](#), GANpaint Studio allows a user to upload an image of their choosing and modify multiple aspects of its appearance, from changing the size of objects to adding completely new items like trees and buildings.

Boon for designers

Spearheaded by MIT professor Antonio Torralba as part of the MIT-IBM Watson AI Lab he directs, the project has vast potential applications. Designers and artists could use it to make quicker tweaks to their visuals. Adapting the system to video clips would enable computer-graphics editors to quickly compose specific arrangements of objects needed for a particular shot. (Imagine, for example, if a director filmed a full scene with actors but forgot to include an object in the background that's important to the plot.)

GANpaint Studio could also be used to improve and debug other GANs that are being developed, by analyzing them for "artifact" units that need to be removed. In a world where opaque AI tools have made image manipulation easier than ever, it could help researchers better understand neural networks and their underlying structures.

"Right now, machine learning systems are these [black boxes](#) that we don't always know how to improve, kind of like those old TV sets that you have to fix by hitting them on the side," says Bau, lead author on a related paper about the system with a team overseen by Torralba. "This research suggests that, while it might be scary to open up the TV and take a look at all the wires, there's going to be a lot of meaningful information in there."

One unexpected discovery is that the system actually seems to have learned some simple rules about the relationships between objects. It somehow knows not to put something somewhere it doesn't belong, like a window in the sky, and it also creates different visuals in different contexts. For example, if there are two different buildings in an image and the system is asked to add doors to both, it doesn't simply add identical doors—they may ultimately look quite different from each other.

"All drawing apps will follow user instructions, but ours might decide not to draw anything if the user commands to put an object in an impossible location," says Torralba. "It's a drawing tool with a strong personality, and it opens a window that allows us to understand how GANs learn to represent the visual world."

GANs are sets of neural networks developed to compete against each other. In this case, one network is a generator focused on creating realistic images, and the second is a discriminator whose goal is to not be fooled by the generator. Every time the discriminator 'catches' the generator, it has to expose the internal reasoning for the decision, which allows the generator to continuously get better.

"It's truly mind-blowing to see how this work enables us to directly see that GANs actually learn something that's beginning to look a bit like common sense," says Jaakko Lehtinen, an associate professor at

Finland's Aalto University who was not involved in the project. "I see this ability as a crucial steppingstone to having autonomous systems that can actually function in the human world, which is infinite, complex and ever-changing."

Stamping out unwanted "fake" images

The team's goal has been to give people more control over GAN networks. But they recognize that with increased power comes the potential for abuse, like using such technologies to doctor photos. Co-author Jun-Yan Zhu says that he believes that better understanding GANs—and the kinds of mistakes they make—will help researchers be able to better stamp out fakery.

"You need to know your opponent before you can defend against it," says Zhu, a postdoc at CSAIL. "This understanding may potentially help us detect fake images more easily."

To develop the system, the team first identified units inside the GAN that correlate with particular types of objects, like trees. It then tested these units individually to see if getting rid of them would cause certain objects to disappear or appear. Importantly, they also identified the units that cause visual errors (artifacts) and worked to remove them to increase the overall quality of the image.

"Whenever GANs generate terribly unrealistic images, the cause of these mistakes has previously been a mystery," says co-author Hendrik Strobelt, a research scientist at IBM. "We found that these mistakes are triggered by specific sets of neurons that we can silence to improve the quality of the image."

Bau, Strobelt, Torralba and Zhu co-wrote the paper with former CSAIL Ph.D. student Bolei Zhou, postdoctoral associate Jonas Wulff, and

undergraduate student William Peebles. They will present it next month at the SIGGRAPH conference in Los Angeles. "This system opens a door into a better understanding of GAN models, and that's going to help us do whatever kind of research we need to do with GANs," says Lehtinen.

This story is republished courtesy of MIT News (web.mit.edu/newsoffice/), a popular site that covers news about MIT research, innovation and teaching.

Provided by Massachusetts Institute of Technology

Citation: Teaching artificial intelligence to create visuals with more common sense (2019, July 1) retrieved 8 December 2023 from <https://techxplore.com/news/2019-07-artificial-intelligence-visuals-common.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.