

TECHNOLOGY

SEE THE DIFFERENCE ONE YEAR MAKES IN ARTIFICIAL INTELLIGENCE RESEARCH


AN IMPROVED WAY OF LEARNING ABOUT NEURAL NETWORKS

Posted May 31,

By Dave Gershgorn 2016

2015


2016



Ant



Starfish



cheeseburger




swimming trunks



Parachute



Screw



running shoe



water jug

Google/ Geometric Intelligence

The difference between Google's generated images of 2015, and the images generated in 2016.

Last June, Google wrote that it was teaching its artificial intelligence algorithms to generate images of objects, or "dream." The A.I. tried to generate pictures of things it had seen before, like dumbbells. But it ran into a few problems. It was able to successfully make objects shaped like dumbbells, but each had disembodied arms sticking out from the handles, because arms and dumbbells were closely associated. Over the course of a year, this process has become incredibly refined, meaning these algorithms are learning much more complete ideas about the world.

New research shows that even when trained on a standardized set of images, A.I. can generate increasingly realistic images of objects that it's seen before. Through this, the researchers were also able to sequence the images and make low-resolution videos of actions like skydiving and playing violin. The paper, from the University of Wyoming, Albert Ludwigs University of Freiburg, and Geometric Intelligence, focuses on deep generator networks, which not only create these images but are able to show how each neuron in the network affects the entire system's understanding.

Looking at generated images from a model is important because it gives researchers a better idea about how their models process data. It's a way to take a look under the hood of algorithms that usually act independent of human intervention as they work. By seeing what computation each neuron in the network does, they can tweak the structure to be faster or more accurate.

"With real images, it is unclear which of their features a neuron has learned," the team wrote. "For example, if a neuron is activated by a picture of a lawn mower on grass, it is unclear if it 'cares about' the grass, but if an image...contains grass, we can be more confident the neuron has learned to pay attention to that context."

They're researching their research—and this gives a valuable tool to continue doing so.



Screenshot

Take a look at some other examples of images the A.I. was able to produce.

ARTIFICIAL INTELLIGENCE , A.I. , DEEP LEARNING , GENERATIVE MODELS , IMAGE PROCESSING , GOOGLE , DEEP DREAM , TECHNOLOGY

YOU MAY LIKE

Sponsored Links by Taboola

Saviez Vous Que Ces 10 Personnalités étaient Gay?

C'est Funny

Ne Parlez Jamais De Ce Jeu A Votre Femme!

League Of Angels II

7 tricks for how to learn a language by yourself

Babbel

The Only 2 Sites You Need to Know About When Building a Website

Top 10 Best Website Builders

EDITORS' PICKS



How Artificial Intelligence Will Translate Facebook Photos For The Blind



Extreme Science



Will The Food Of The Future Be 'Unnaturally Delicious'?

LATEST NEWS

Bill Nye's Solar Sail Hits A Few Snags But Is Almost Ready To Fly

Watch Jimmy Kimmel Become Mike Tyson And Karl Malone Digitally

Google Has A List Of A.I. Behaviors That Would Scare Them Most

Different Scorpion Species Share Similar Taste In Interior Design

Smelly Postcards Will Give You A Whiff Of Rosetta's Comet

VIDEOS



The Sonic Science Of 'Hamilton'



Watch Nvidia's Autonomous Car Drive Through Snow And Winding Roads



Q Site Search

[SUBSCRIBE](#) [RENEW](#) [CUSTOMER SERVICE](#) [PRIVACY POLICY](#) [TERMS OF USE](#) [MASTHEAD](#) [CONTACT US](#) [POPULAR SCIENCE TV](#)

Copyright © 2016 Popular Science. A [Bonnier Corporation](#) Company. All rights reserved. Reproduction in whole or in part without permission is prohibited.

BONNIER
Corporation

Visit Our Sister Sites 