


ARTIFICIAL INTELLIGENCE

Astronomers explore uses for AI-generated images

Neural networks create pictures to train image-recognition programs and science software.

LEFT FIGURE: S. RAVANBAKHSH/DATA. ARXIV.ORG/ABS/1609.05796;
RIGHT: NGUYEN ET AL./ARXIV.ORG/ABS/1612.00005



AI-generated images of galaxies (left, lower of each pair) and volcanoes.

BY DAVIDE CASTELVECCHI

Volcanoes, monasteries, birds, thistles: the varied images in Jeff Clune's research paper could be his holiday snaps. In fact, the pictures are synthetic. They are generated by deep-learning neural networks: layers of computational units that mimic how neurons are connected in the brain.

In recent years, neural networks have made huge strides in learning to recognize and interpret information in pictures, videos and speech. But now, computer scientists such as Clune are turning those artificial-intelligence (AI) systems on their heads to create 'generative' networks that churn out their own realistic-seeming information. "I have reached a point in my life where I'm kind of having reality vertigo," says Clune, who works at the University of Wyoming in Laramie.

Generative systems also give an insight into how neural networks interpret the world, says Kyle Cranmer, a particle physicist and computer scientist at New York University. Although it's not clear how virtual neurons store and interpret information,

the plausibility of the data they generate suggests that they have some handle on the real world.

AI researchers are excited about using generative networks to train image-recognition software. More widely, Cranmer says, AIs that generate scientific data might help astronomers and other researchers to prune noise from large data sets, and so better understand patterns within them.

AI DUEL

In computer science, there's a particular buzz around a technique that sets a generative network in competition with an image-recognition network, to help both to improve their performance. These 'generative adversarial nets', or GANs, are "the coolest idea in deep learning in the last 20 years", said Yann LeCun — head of Facebook's AI team in New York City — in a talk at Carnegie Mellon University in Pittsburgh, Pennsylvania, last November.

Typically, a neural network learns to discriminate between training images tagged by people, with descriptions such as 'Victorian house' or 'golden retriever'. The training

process tells the AI how to tweak connections between its virtual neurons, so that it can eventually tag images by itself — including new photos that were not part of the original training set.

In a GAN, two neural networks train together with minimal outside help. One network, the generator, produces fake images; the other, the discriminator, tries to tell those fake images from real ones. After that, the discriminator checks which images were real and which were fake, so that it can get better at distinguishing between them. The generator never sees the real images — instead, the discriminator tells it how to tweak its output to make its pictures more like the real thing.

"You can think of the discriminator as a teacher that tells the generator how to improve," says Ian Goodfellow, a computer scientist at the non-profit organization OpenAI in San Francisco, California. Or, to put it another way, the discriminator is like a banker who helps a counterfeiter learn how to forge money, he adds. Goodfellow came up with the idea of GANs in 2014, while he was a student of machine-learning pioneer Yoshua Bengio at the University of Montreal in Canada. A game-theory analysis, he says, shows that, in principle, the generator will eventually become so good that the discriminator can no longer tell the difference between real and fake.

AI image-recognition systems can learn much more efficiently using GANs than can conventional deep-learning systems, says Goodfellow: they can become proficient on the basis of hundreds of training images, whereas current state-of-the-art image recognition typically requires tens of thousands. He says that might help in applications such as medical diagnostics, where large sets of patient data exist but are mostly off limits owing to privacy concerns.

AI researchers have invented a variety of approaches to generating images. One, called a variational autoencoder (VAE), seems to be able to produce slightly less realistic but more diverse images than GANs, says computer scientist Max Welling at the University of Amsterdam. And other teams have developed further variants, some combining GANs and VAEs. Clune, Bengio and others collaborated on combined networks to generate the

photo-realistic images in their latest paper, which was posted on the arXiv preprint server last November (A. Nguyen *et al.* Preprint at <http://arxiv.org/abs/1612.00005>; 2016).

ARTIFICIAL DATA

Generative AIs look promising for basic science, too, says Welling, who is helping to develop software for the Square Kilometre Array (SKA), a radio-astronomy observatory to be built in South Africa and Australia. The SKA will produce such vast amounts of data that its images will need to be compressed into low-noise but patchy data. Generative AI models will help to reconstruct and fill in blank parts of those data, producing the images of the sky that astronomers will examine.

A team led by Rachel Mandelbaum, an astrophysicist at Carnegie Mellon University, has been experimenting with both GANs and VAEs to simulate images of galaxies that look deformed because of gravitational lensing — when the gravity of objects in the foreground distorts space-time and warps light rays. Researchers are planning to survey huge numbers of galaxies to map gravitational lensing across the Universe's history. This could show how the distribution of the Universe's matter has changed over time, providing clues to the nature of the dark energy that is thought to have driven cosmic expansion. But to do this, astronomers need software that can reliably separate gravitational lensing from other effects. Synthetic images will improve the programs' accuracy, Mandelbaum says.

Many scientists hope that the latest AI neural nets will help them to discover patterns in huge, complex sets of data — but some are wary of trusting the interpretation of such 'black box' systems, whose inner workings are mysterious. Even if virtual neurons seem to give correct answers, they might have a mistaken understanding of the world. But adding a generative element to a neural net could help, Cranmer says. "If it can generate data that look just as real, then it's much more convincing that, whatever the black box is, it has actually learned the physics."

Clune worries about generative algorithms, too. For all their potential benefits, he's concerned about the social implications of having machines that will one day be able to produce fake but real-looking pictures or video — perhaps, say, of Donald Trump receiving bribes from Vladimir Putin. "I think that, increasingly, this is going to be an interesting challenge in society," he says. ■

PHYSICS

Hydrogen yet to prove it's metal

Doubts shroud claimed fulfilment of 80-year quest.

BY DAVIDE CASTELVECCHI

Two physicists say that they have crushed hydrogen under such immense pressures that the gas became a shiny metal — a feat that scientists have been trying to accomplish for more than 80 years.

But other researchers have serious doubts about the claim, the latest in a field with a long history of failed attempts. Ranga Dias and Isaac Silvera, both physicists at Harvard University in Cambridge, Massachusetts, first posted a report of their results on the arXiv preprint server last October¹, which attracted immediate criticism. A peer-reviewed version was published on 26 January in *Science*², but sceptics say that it includes little new information.

Five experts told *Nature's* news team that they do not yet believe the claim, and need more evidence. "I don't think the paper is convincing at all," says Paul Loubeyre, a physicist at France's Atomic Energy Commission in Bruyères-le-Châtel.

Silvera and Dias say that they wanted to publish their first observation before making further tests on their fragile material.

METALLIC DREAM

Creating metallic hydrogen has been a dream in the high-pressure-physics world since 1935, when theorists predicted its existence³. Squeezed hard enough inside an anvil, hydrogen should be able to conduct electricity, the hallmark of a metallic state. Theorists say that the material could have other exotic properties, such as being a superconductor — able to conduct electricity without resistance — at room temperature.

By making metallic hydrogen, physicists might also be able to explore planetary science at their lab bench: gas-giant planets such as Jupiter are theorized to have metallic hydrogen in their cores, which would perhaps explain how they can sustain a magnetic field.

In recent years, physicists have crushed tiny samples of hydrogen between diamond anvils at pressures exceeding those at the centre of Earth. But no one has made metallic hydrogen. Dias and Silvera say that they were able to squeeze their hydrogen gas at pressures greater than those achieved by anyone else. To do so, they used an anvil that can fit inside a cryostat, enabling them to cool the hydrogen

to just above absolute zero. They also say that they have found a better way to polish the tips of their diamonds, to remove irregularities that could break the gems. They then turned a screw to crank up the pressure to 495 billion pascals, or almost 5 million times higher than atmospheric pressure at sea level.

"Then, suddenly, it becomes a lustrous, reflective sample, which you can only believe is a metal," Silvera says. Seen through a microscope, the sample appeared shiny, and it reflected light in the way metallic hydrogen should do, he says.

REDO THE MEASUREMENT

It's far from clear, however, that the shiny material the researchers see is actually hydrogen, says Alexander Goncharov, a geophysicist at the Carnegie Institution for Science in Washington DC. He suggests that the material could be alumina (aluminium oxide), which coats the tips of the diamonds in the anvil, and might behave differently under pressure.

Loubeyre and others think that Silvera and Dias are overestimating the pressure that they produced, by relying on an imprecise calibration between turns of the screw and pressure inside the anvil. Eugene Gregoryanz, a physicist at the University of Edinburgh, UK, adds that part of the problem is that the researchers took only a single detailed measurement of their sample at the highest pressure — making it hard to see how pressure shifted during the experiment.

"If they want to be convincing, they have to redo the measurement, really measuring the evolution of pressure," says Loubeyre. "Then they have to show that, in this pressure range, the alumina is not becoming metallic."

But Silvera says that he and Dias just wanted to get the news out before making confirmation tests, which could break their precious specimen. "We wanted to publish this breakthrough event on this sample," he says. To preserve the material, they have kept it in the cryostat; the lab has only two cryostats, and the other is in use for other experiments. "Now that the paper has been accepted," he says, "we're going to do further experiments. ■"

1. Dias, R. & Silvera, I. F. Preprint at <https://arxiv.org/abs/1610.01634> (2016).
2. Dias, R. P. & Silvera, I. F. *Science* <http://dx.doi.org/10.1126/science.aal1579> (2017).
3. Wigner, E. & Huntington, H. B. *J. Chem. Phys.* **3**, 764–770 (1935).