

Machine Learning Optimizes Images for Stimulating Monkey Neurons

7-9 minutes

Showing monkeys a series of computer-generated images and simultaneously recording the animals' brain cell activities enables deep machine learning systems to generate new images that ramp up the cells' excitation, according to two papers published today (May 2) in [Cell](#) and [Science](#).

"It's exciting because it's bridging the fields of deep learning and neuroscience . . . to try and understand what is represented in different parts of the brain," says neuroscientist [Andreas Tolias](#) of Baylor College of Medicine who was not involved with either of the studies, but has carried out [similar experiments](#) in mice. "I think these methods and their further development could provide a more systematic way for us to open the black box of the brain," he says.

It's a goal of sensory neuroscience to understand exactly which stimuli activate which brain cells. In the primate visual system, certain neurons in the visual cortex and inferior temporal cortex (two key vision areas) are known to respond preferentially to certain stimuli—such as colors, specific directions of motion, curves, and even faces. But, says neuroscientist [Carlos Ponce](#) of Washington University School of Medicine in St. Louis, who co-authored the [Cell](#) paper, "the problem is, we've never quite known whether, in our selection of pictures, we have the secret true image that the cell really is encoding." Maybe, he suggests, a cell isn't responding to a face, but to an arrangement of features and shapes found in a face that may also be found in other images. And with countless available images, "it's impossible to test all of them," he says. In short, it has been impossible to determine the exact visual stimulus that would maximally activate a given neuron.

"We've been stuck like this for the past 50 years, unable to solve this problem," says Ponce. Recent advances in machine learning, specifically the development of generative adversarial networks (GANs), have now enabled a solution.

GANs are neural networks that can, says Ponce, "take the visual world and reproduce it in their own internal abstractions." That is, they can generate an infinite variety of realistic and abstract pictures based on example images fed to them.

"We thought, why don't we take this amazing network, which can simulate the world, and

connect it somehow to neurons in an actual monkey's brain," says Ponce. He explains that by feeding back information to the network about which images activate a neuron most potently, the network could tweak the images to boost activity further—thus finding that elusive “secret true image” the cell encodes.



An image evolved by a neural network to stimulate a macaque neuron

PONCE, XIAO AND SCHADE ET AL.

To do this, Ponce's team implanted microelectrode arrays in the inferior temporal cortices of six monkeys, an area slightly above and behind the ears, to record the activities of vision neurons. They then showed the animals 40 random images generated by a GAN and identified the 10 images that induced the most activity in a given neuron or group of neurons. The researchers put those images through a genetic algorithm—essentially recombining the pixels—to create 30 similar images, and then showed those images, together with the original top 10, to the monkey. This cycle was repeated up to 250 times over one to three hours until the animal was no longer looking at the images. As controls, the team also included naturalistic images of people, places, and simple line drawings.

The images that resulted from this evolutionary process were a mix of objects, shapes and colors, some of which exhibited a partial resemblance to monkeys, or the faces of the animals' handlers. Importantly, when such image evolution was repeated for the same neuron in the same animal on a different day, a similar image emerged.

“To see from this mist a face-like object stare back at you, when you know this image is coming from a monkey's brain, is probably one of the most magical experiences I've had in science,” says Ponce.

As the images evolved, they more strongly stimulated the cells guiding the evolution, the

team found. “The [evolved images] drive these neurons two times better, or three times better” than the original images or natural images that stimulated the neurons, “so it’s a huge effect,” says neuroscientist [Christof Koch](#) of the Allen Institute for Brain Science who was not involved in either project. “It’s really surprising that it worked so damn well.”

Indeed “both of these papers show that these techniques drive neurons in much higher capacities than previous standard images,” says [Kohitij Kar](#), a neuroscience postdoc in the lab of [James DiCarlo](#) at MIT, referring to Ponce’s study and his own in [Science](#).

Kar and his colleagues’ approach was also to show neural network-generated images to monkeys with implanted microelectrode arrays. In this team’s case, however, the aim was to test the strength of the network as a model for the visual system.

While neural networks are generally considered a good model of the brain’s visual system—indeed, their architecture is based on neural circuitry—there has been limited testing of their functional similarity.

To see from this mist a face-like object stare back at you, when you know this image is coming from a monkey’s brain, is probably one of the most magical experiences I’ve had in science.

—Carlos Ponce, Washington University School of Medicine in St. Louis

To remedy this, Kar and colleagues showed the same set of images to a monkey and to a neural network and correlated the particular neurons that fired in response to a given image with the output of particular units (mathematical functions) in the network. Using this correlative information, the network made new images that, based on its own mathematical functions, were predicted to boost the activity of corresponding neurons in the animal’s brain.

“If these units in the artificial neural networks and [the neurons] in the brain are doing the same thing then they should respond to similar images in the same way,” explains [Pouya Bashivan](#), a coauthor on the *Science* paper who is also a postdoc in DiCarlo’s lab.

Sure enough, “we showed that about 70 percent of the time, by following this procedure we can stretch the neural responses higher than what we had seen in response to natural images,” says Bashivan. The team expects to increase that rate in the future by further improving the models.

“I think we are heading into a future where. . . neural networks one day will be excellent models of the brain,” says Ponce, “but we need to get more information about how the actual brain works to try and inform the development of future neural networks.”

It’s a circular process, says Koch. “The brain was an inspiration for neural networks. Neural networks are now immensely more powerful than they were 50 years ago, and in fact they are rivaling us . . . not only that, they are also helping us understand the vastly more complex networks that we all have in our heads.”

C.R. Ponce et al., “Evolving images for visual neurons using a deep generative network reveals coding principles and neuronal preferences,” [Cell](#), doi: 10.1016/j.cell.2019.04.005, 2019.

P. Bashivan et al., “Neural population control via deep image synthesis,” [Science](#), 364, eaav9436, 2019.