

Nonlinear optical computing doesn't need nonlinear optics

A major stumbling block on the road to light-based neural networks can be overcome by flipping the script on how data are encoded.

Artificial neural networks are making their mark on both the world and its energy budget. The brain-inspired computing models behind so many

popular and scientific machine-learning applications are proving to be tremendously powerful (see, for example, *PHYSICS TODAY*, October 2021, page 14). But they're also power hungry (see *PHYSICS TODAY*, April 2024, page 28).

One way, potentially, to lessen the energy burden is to design a computer that uses light, not electrons, to process data. Most of the computations that a neural network needs to do are linear operations: adding, subtracting,

and multiplying by constants. An optical computer could perform those operations quickly and energy efficiently.

As adept as optical computing is with linear computations, though, it struggles mightily with nonlinear ones—a small but necessary ingredient in neural-network computing—for the simple reason that photons don't generally interact with one another. Some nonlinear optical materials can mediate




FIGURE 1. INSPIRED BY THE BRAIN, artificial neural networks process information by passing it among layers of nodes, known as neurons. Along the way are matrices of trainable parameters (represented in the inset as J_1 and J_2) which are iteratively adjusted to optimize the network to perform a specific task. Here, that task is recognizing images of numerals. (Image courtesy of Clara Wanjura; inset adapted from ref. 1.)




FIGURE 2. TO USE LIGHT as the basis for a neural network, researchers must rethink how they encode and process data. **(a)** A spherical cavity, partially lined with a reconfigurable array of mirrors, produces a random-looking speckle pattern when light bounces around inside it. But the speckles carry some surprisingly detailed information about an image encoded in the mirror array. (Courtesy of Fei Xia.) **(b)** In a more programmable implementation, four copies of an input image are encoded in a spatial light modulator, and an illumination beam is scattered off all of them. The system can be trained to classify the image, even one that it's never seen before. (Adapted from ref. 3.)

light–light interactions and thereby produce nonlinear responses, but they typically require impractically high optical power.

That chain of reasoning assumes that data are encoded in a light field, which is then processed and manipulated by the optical neural network. But as three groups have now shown, that's not the only way to do it. The groups' approaches differ, but their common insight is to encode the input data not in the light itself but in some part of the system with which the light interacts. Nonlinear functions can then be computed with ease, and neural-network implementations follow.

At the Max Planck Institute for the Science of Light in Erlangen, Germany, Clara Wanjura and Florian Marquardt showed theoretically that input data can be represented as frequency offsets in a system of coupled resonators.¹ The other two groups—one led by Hui Cao of Yale University and Sylvain Gigan of the École Normale Supérieure in Paris,² and the other led by Demetri Psaltis and Christophe Moser at the Swiss Federal Institute of Technology in Lausanne (EPFL)³—experimentally encoded data in arrays of pixels that light scatters off multiple times. In each case, the light-based systems can tackle rudimentary image-classifica-

tion tasks with accuracies on par with digital neural networks.

Enlightened data processing

A neural network is really just a fancy mathematical function. It takes an input, such as the grainy image of a handwritten numeral shown in the inset in figure 1, and it produces an output: "7". To get from one to the other, it processes the data through layers of nodes, or neurons.

In a conventional neural network, each neuron computes a weighted average of all the signals feeding into it. Then, depending on whether the result exceeds a certain threshold, the neuron either fires or doesn't fire—that is, it produces either a 1 or a 0 to feed into the next layer. The weights used in the weighted average are so-called trainable parameters: The model is iteratively adjusted based on a series of inputs whose correct outputs are known, until eventually the network can correctly process inputs that it's never seen before.

The weighted averages, which consume the bulk of the computing power, even during the training phase, are the kind of linear operations at which optical computing excels. Summing two optical signals is as simple as superposing two light fields. And

even a more complicated series of weighted sums can be done easily with a network of beamsplitters and phase shifters.⁴

For linear operations, optical computing can outshine electronic computing for all the same reasons as fiber-optic cables excel at transmitting data over long distances: Information encoded in an optical beam can be tightly compressed in both space and time, and it can travel long distances with little dissipation. So data can be processed with high throughput and low power.

But the big hurdle for optical neural networks is the simplest-sounding part of the computation: the decision of each neuron to fire or not, which is a nonlinear function of the input data. It can be computed with nonlinear optics, albeit with high optical power, or by converting signals from optical to electronic and back. But those approaches negate some of optical computing's biggest advantages.

Happily, neural networks aren't too picky about the nature of the nonlinear function. It doesn't have to be an all-or-nothing step function. In fact, most implementations use a smoothed step function for ease of computation, and many other nonlinear functions can be made to work if the network is suitably trained. So the question becomes, can a

platform of all linear optics re-create any nonlinear function at all?

House of mirrors

The answer to that question—a resounding yes—depends on asking the right follow-up question: A nonlinear function of what? Linear optics, by definition, can compute only linear functions of an input light field, but they can produce a nonlinear response to other physical parameters, such as the positions of mirrors. To use those nonlinearities as the basis for a neural network, though, requires making some big changes to how the network is structured.

Cao and her group at Yale came to the study of optical nonlinearities while pursuing a different application: Rather than a neural network, they wanted to create a physical unclonable function (PUF), a sort of digital fingerprint that can serve as a security feature in the internet of things.⁵ They took a golf ball-sized spherical cavity, as shown in figure 2a, and lined part of the interior with a reconfigurable array of tiny mirrors and the rest with a diffuse reflective coating. When they shot a laser beam into a hole in the cavity, the light bounced around inside before emerging as a speckle pattern out another hole.

The output speckles depend deterministically and reproducibly on the mirror configuration, but in a way that's almost impossible to replicate by anyone not in possession of that specific cavity. Those properties are what make the system a PUF—but is it also a neural network? It wouldn't initially seem to be: It has no discernible neurons, weighted averages, or trainable parameters. But in collaboration with Gigan and his postdoc Fei Xia, Cao and colleagues realized that the cavity could act as a so-called reservoir computer, a type of neural network in which all the computing is done first and all the interpretation is done later.

The speckle pattern is a highly nonlinear function of the mirror configuration. On average, the researchers estimate, light bounces thousands of times off the cavity surface before exiting, and at least a few hundred of those bounces are off the mirror array. The resulting speckles are full of information about correlations among pixels in the input data. And correlations are just what all neural networks use to work their magic.

To make sense of the speckle pattern, the researchers need only pass it through a decoder with one or a few layers of trainable weights, which is easy enough to do electronically. The mirror array, with more than 4 million pixels, can encode extremely detailed input images, and the system is capable of some complex computing tasks, including recognizing subtle features of human faces and spotting pedestrians in traffic scenes. Those tasks require up to 1 million trainable parameters in the output decoder. But that's less than a conventional neural network uses for the same tasks.

The EPFL researchers were also inspired by Cao and colleagues' PUF paper, but they took their implementation in a different direction. "We wanted to maintain a degree of programmability," says Mustafa Yildirim, one of the paper's co-first authors along with Niyazi Ulas Dinç. Like Cao, Gigan, and colleagues, they generate their nonlinearity by bouncing a light beam off an input image multiple times. But instead of leaving the scattering dynamics up to random chance, they send the light on a controlled zigzag path that scatters off a spatial light modulator with four distinct copies of the input, as shown in figure 2b.

The four copies of the input aren't all identical. In each one, every pixel is linearly scaled by a pair of trainable parameters, so the researchers can train their network much like a conventional one. Although the light bounces only four times off the input, the degree of nonlinearity was sufficient for the EPFL researchers to successfully train their system to perform simple image classifications, like distinguishing pictures of dogs, fish, and T-shirts. And because the input data are encoded four separate times, the network is highly robust against noise.

New architectures

Wanjura and Marquardt's work is the most abstract of the three groups. As theorists, they're focused on the mathematical concepts behind neuromorphic computing schemes. "I'd previously employed scattering theory in my topology studies," says Wanjura. "When I read Florian Marquardt's lecture notes on machine learning, I noticed that the scattering matrices I had worked with had

some similarity to the math behind neural networks. So when I joined his group as a postdoc, we developed the idea further together."

Like a conventional neural network, the one that Wanjura and Marquardt envisioned is made up of discrete neurons. But unlike a conventional neural network, the information in it doesn't flow only one way. Rather, the light waves—or any waves, really—scatter back and forth across the network in both directions. With the input data and trainable parameters encoded in some of the neurons, the optical signal picks up a nonlinear dependence on both.

Wanjura and Marquardt proposed that the network could be realized by a system of coupled resonators, in which information is encoded in a resonator by detuning it from resonance. They're collaborating with Amir Safavi-Naeini and his experimental group at Stanford University to bring their ideas to fruition. But as a first step, they ran simulations of their network on an ordinary computer to show that it works for classifying images of handwritten numerals. "It's ironic," says Wanjura, "that the training simulations on a computer required a few hours, whereas a photonics experiment could ideally perform the entire training in a few milliseconds."

All three of the groups' endeavors are still at the proof-of-principle stage. Because all their networks process data so differently from conventional neural networks, it remains to be seen whether any of them can be scaled up to rival the powerful, power-hungry hardware that operates applications such as ChatGPT. But the implementations show that there's potential value in thinking outside the box. "It's motivating us to take more risks in optical computing," says Dinç, "not just directly adapting everything from electronics."

Johanna Miller

References

1. C. C. Wanjura, F. Marquardt, *Nat. Phys.* (2024), doi:10.1038/s41567-024-02534-9.
2. F. Xia et al., *Nat. Photon.* (2024), doi:10.1038/s41566-024-01493-0.
3. M. Yildirim et al., *Nat. Photon.* (2024), doi:10.1038/s41566-024-01494-z.
4. D. A. B. Miller, *Photon. Res.* **1**, 1 (2013).
5. Y. Eliezer et al., *Proc. Natl. Acad. Sci. USA* **120**, e2305027120 (2023).