TUTORIAL

A Photonic Neuron

The light-driven building block at the heart of computers for next-gen machine learning.

Artificial neural networks (ANNs) driven by software algorithms are enabling many new applications. Yet software-driven ANNs are a poor fit with existing computer architectures. For that reason, researchers are hard at work on "neuromorphic" computers, in which the hardware itself is inspired by the networks of the human brain.

At the heart of such new computers will be networks of interlinked artificial neurons—and integrated photonics will play a key role in building and wiring them up. To learn about what goes into a "photonic neuron," OPN talked with Bhavin Shastri of Queen's University, Ontario, Canada.

1. From biology to computation

At its simplest level, a biological neuron takes in signals from other neurons, integrates them biochemically, and, if the integrated signal is strong enough, fires off its own pulse to still other neurons. This is the "integrate and fire" model that underlies many ANNs.

In such a model, Shastri explains, an artificial neuron has two parts: A linear part, in which inputs from many other neurons are weighted and summed together; and a nonlinear part, in which the artificial neuron "fires" if the inputs exceed a given threshold. "The linear part is essentially a linear filter, doing a pattern correlation," he says, "and the nonlinear activation function is doing the decision-making" on whether the match is strong enough.

2. Weighting the inputs

Weighting the inputs from many neurons, Shastri observes, is accomplished via matrix multiplication. "You are doing a dot product of two vectors," he says—the array of input values from other neurons, and the array of weights applied to them. And those weights need to be tunable, as adjusting them is how the neural network is trained.

One way to implement this photonically, which Shastri and colleagues have demonstrated, is to use wavelength-division multiplexing (WDM) to "fan in" optical inputs from other neurons as different-wavelength signals in the same waveguide. Then, tunable ring resonators can be used to attenuate intensities at different wavelengths, thereby applying the weighting. Shastri stresses that there are other ways to do this kind of weighting photonically, too—some teams, for example, have used meshes of Mach-Zehnder interferometers to implement weighting.

This linear weighting function is where photonics really shines over all-electronic approaches. With electronics, the many matrix multiplications required impose a significant energy cost. With photonics, Shastri says, "it's happening in the time of flight," with far lower energy input. "And thousands of signals can coexist in a single waveguide."

3. Summing up

In the system shown here, the next step—integrating the weighted input signals—is done by a balanced photodiode that sums the optical signals at many wavelengths into a single photocurrent. Some groups, Shastri says, are also attempting to use all-optical approaches, avoiding electronics entirely. The right mixture of electronics and optics in these systems is still a point of active debate.

4. The neuron fires

Once the input signals are weighted and summed, the neuron must decide whether they're strong enough to propagate further in the network. This is the nonlinear activation function—and how it's implemented depends in part on the kind of neural network.

In one model—a "spiking" neural network, shown here—the nonlinear step could be executed by a combination of a laser cavity, an electrically pumped gain medium and a saturable absorber. Once again, though, Shastri stresses that research groups are looking at many possible approaches. For example, the use of electro-optic modulators to handle nonlinear activation in "non-spiking" neural networks is advancing rapidly—and is likely a much better route to large-scale integration in silicon.

5. One of many

A real ANN, of course, includes multiple photonic neurons, each wired up via optical waveguides to multiple others. Getting to practical neuromorphic computers will require significant work packing these neurons—as well as required I/O systems, analog memories, light sources and digital components—into a CMOS-compatible platform. It's a challenge on which Shastri and many others in the community are working now. OPN

OPN thanks Bhavin Shastri and Alexander Tait for assistance with this tutorial. For references and resources, go online: www.osa-opn.org/tutorial/photonic-neuron.

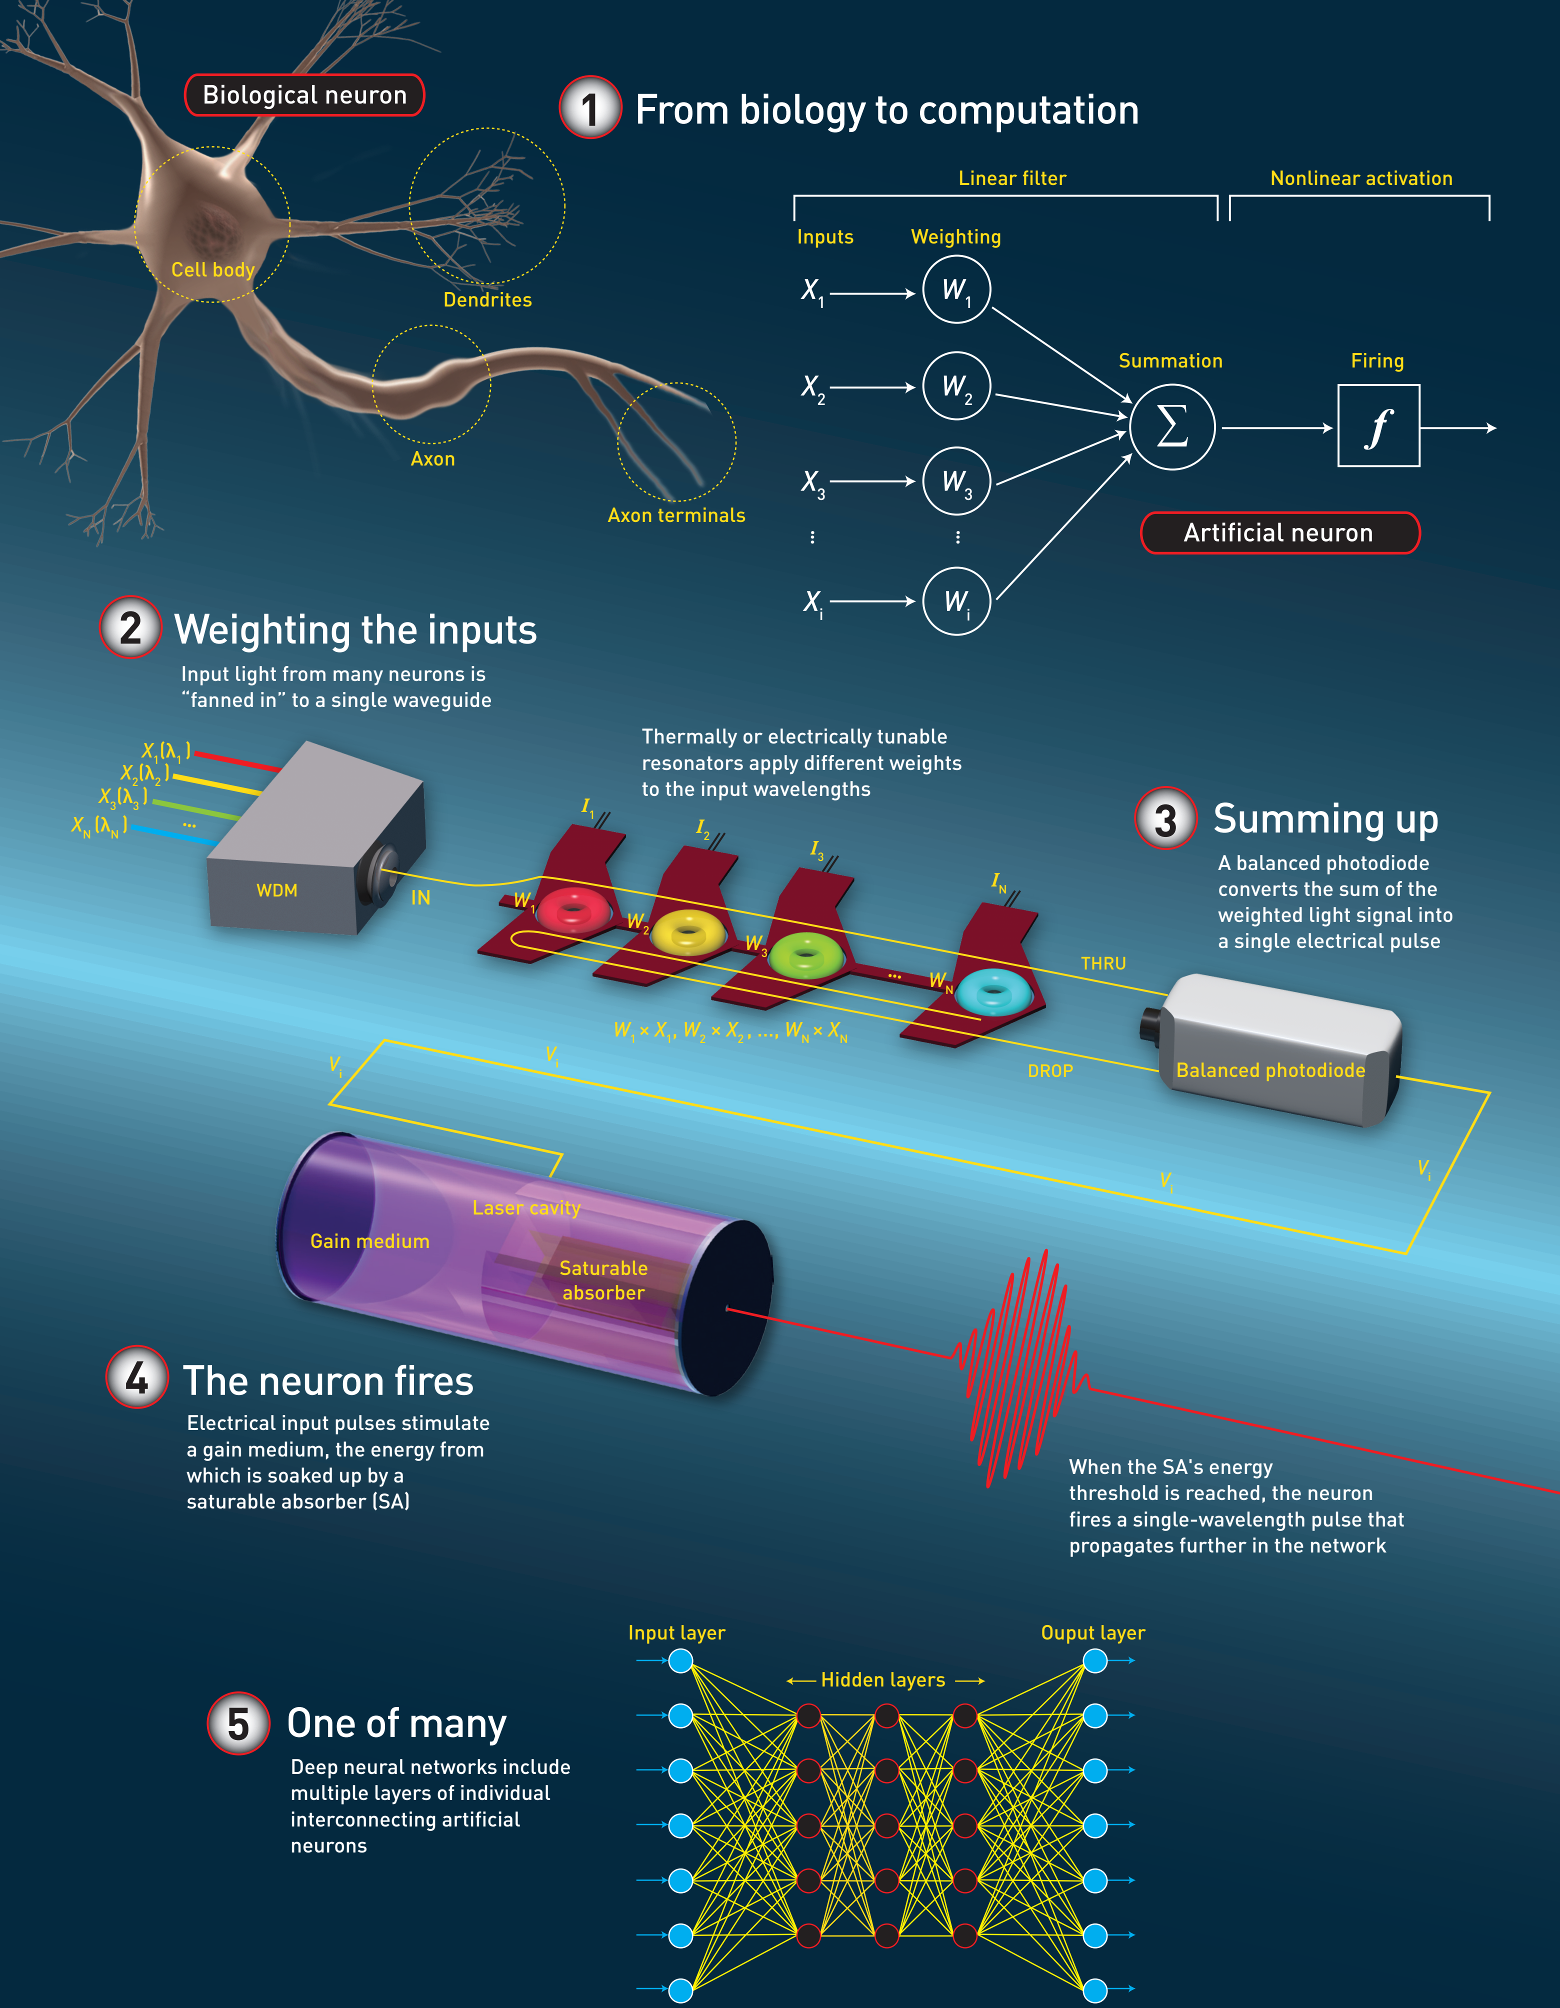

Illustration by Phil Saunders