While modern artificial neural networks are the backbone of the current AI revolution, they are only loosely inspired by networks of real, biological neurons such as our brain. The brain however is a much larger network, much more energy-efficient, and can respond ultra-fast when triggered by external events. Spiking neural networks are special types of neural networks that more closely mimic the working of biological neurons: the neurons of our nervous system communicate by exchanging electrical pulses, and they do so only sparingly.

Implemented in chips, called neuromorphic hardware, such spiking neural networks hold the promise of bringing AI programmes closer to users – on their own devices. These local solutions are good for privacy, robustness and responsiveness. Applications range from speech recognition in toys and appliances, health care monitoring and drone navigation to local surveillance.

Just like standard artificial neural networks, spiking neural networks need to be trained to perform such tasks well. However, the way in which such networks communicate poses serious challenges. "The algorithms needed for this require a lot of computer memory, allowing us to only train small network models mostly for smaller tasks. This holds back many practical AI applications so far," says Sander Bohté of CWI's Machine Learning group. In the Human Brain Project, he works on architectures and learning methods for hierarchical cognitive processing.

Mimicking the learning brain

The learning aspect of these algorithms is a big challenge, and they cannot match the learning ability of our brain. The brain can easily learn immediately from new experiences, by changing connections, or even by making new ones. The brain also needs far fewer examples to learn something and it works more energy-efficiently. "We wanted to develop something closer to the way our brain learns," says Bojian Yin.

Yin explains how this works: if you make a mistake during a driving lesson, you learn from it immediately. You correct your behaviour right away and not an hour later. "You learn, as it were, while taking in the new information. We wanted to mimic that by giving each neuron of the neural network a bit of information that is constantly updated. That way, the network learns how the information changes and doesn't have to remember all the previous information. This is the big difference from current networks, which have to work with all the previous changes. The current way of learning requires enormous computing power and thus a lot of memory and energy."

Six million neurons

The new online learning algorithm makes it possible to learn directly from the data, enabling much larger spiking neural networks. Together with researchers from TU Eindhoven and research partner Holst Centre, Bohté and Yin demonstrated this in a system designed for recognising and locating objects. Yin shows a video of a busy street in Amsterdam: the underlying spiking neural network, SPYv4, has been trained in such a way that it can distinguish cyclists, pedestrians and cars and indicate exactly where they are.

"Previously, we could train neural networks with up to 10,000 neurons; now, we can do the same quite easily for networks with more than six million neurons," says Bohté. "With this, we can train highly capable spiking neural networks like our ¬¬SPYv4."

Future

And where does it all lead? Now having such powerful AI solutions based on spiking neural networks, chips are being developed that can run these AI programmes at very low power and will ultimately show up in many smart devices, like hearing-aides and augmented or virtual reality glasses.

Original Publication:

The researchers