Spiking tool improves artificially intelligent devices

February 27 2019, by Neal Singer

Against a background of more conventional technologies, Sandia National Laboratories researchers, from left, Steve Verzi, William Severa, Brad Aimone and Craig Vineyard hold different versions of emerging neuromorphic hardware platforms. The Whetstone approach makes artificial intelligence algorithms more efficient, enabling them to be implemented on smaller, less power-hungry hardware. Credit: Randy Montoya

Whetstone, a software tool that sharpens the output of artificial neurons, has enabled neural computer networks to process information up to a hundred times more efficiently than the current industry standard, say the Sandia National Laboratories researchers who developed it.

The aptly named software, which greatly reduces the amount of circuitry needed to perform autonomous tasks, is expected to increase the penetration of artificial intelligence into markets for mobile phones, self-driving cars and automated interpretation of images.

"Instead of sending out endless energy dribbles of information," Sandia neuroscientist Brad Aimone said, "artificial neurons trained by Whetstone release energy in spikes, much like human neurons do."

The largest artificial intelligence companies have produced spiking tools for their own products, but none are as fast or efficient as Whetstone, says Sandia mathematician William Severa. "Large companies are aware of this process and have built similar systems, but often theirs work only for their own designs. Whetstone will work on many neural platforms."

The open-source code was recently featured in a technical article in Nature Machine Intelligence and has been proposed by Sandia for a patent.

How to sharpen neurons

Artificial neurons are basically capacitors that absorb and sum electrical charges they then release in tiny bursts of electricity. Computer chips, termed "neuromorphic systems," assemble neural networks into large groupings that mimic the human brain by sending electrical stimuli to neurons firing in no predictable order. This contrasts with a more lock-step procedure used by desktop computers with their pre-set electronic processes.

Because of their haphazard firing, neuromorphic systems often are slower than conventional computers but also require far less energy to operate. They also require a different approach to programming because otherwise their artificial neurons fire too often or not often enough, which has been a problem in bringing them online commercially.

Whetstone, which functions as a supplemental computer code tacked on to more conventional software training programs, trains and sharpens artificial neurons by leveraging those that spike only when a sufficient amount of energy—read, information —has been collected. The training has proved effective in improving standard neural networks and is in process of being evaluated for the emerging technology of neuromorphic systems.

Catherine Schuman, a neural network researcher at Oak Ridge National Laboratories, said, "Whetstone is an important tool for the neuromorphic community. It provides a standardized way to train traditional neural networks that are amenable for deployment on neuromorphic systems, which had previously been done in an ad hoc manner."

The strict teacher

The Whetstone process, Aimone said, can be visualized as controlling a class of talkative elementary school students who are tasked with identifying an object on their teacher's desk. Prior to Whetstone, the students sent a continuous stream of sensor input to their formerly overwhelmed teacher, who had to listen to all of it—every bump and giggle, so to speak—before passing a decision into the neural system. This huge amount of information often requires cloud-based computation to process, or the addition of more local computing equipment combined with a sharp increase in electrical power. Both options increase the time and cost of commercial artificial intelligence products, lessen their security and privacy and make their acceptance

less likely.

Under Whetstone, their newly strict teacher only pays attention to a simple "yes" or "no" measurement of each student—when they raise their hands with a solution, rather than to everything they are saying. Suppose, for example, the intent is to identify whether a piece of green fruit on the teacher's desk is an apple. Each student is a sensor that may respond to a different quality of what may be an apple: Does it have the correct quality of smell, taste, texture and so on? And while the student who looks for red may vote "no" the other student who looks for green would vote "yes." When the number of answers, either yay or nay, is electrically high enough to trigger the neuron's capacity to fire, that simple result, instead of endless waffling, enters the overall neural system.

While Whetstone simplifications could potentially increase errors, the overwhelming number of participating neurons—often over a million—provide information that statistically make up for the inaccuracies introduced by the data simplification, Severa said, responsible for the mathematics of the program.

"Combining overly detailed internal information with the huge number of neurons reporting in is a kind of double booking," he says. "It's unnecessary. Our results tell us the classical way—calculating everything without simplifying—is wasteful. That is why we can save energy and do it well."

Patched programs work best

The software program works best when patched in to programs meant to train new artificial-intelligence equipment, so Whetstone doesn't have to overcome learned patterns with already established energy minimums.

Provided by Sandia National Laboratories