accenture

NEUROMORPHIC COMPUTING: ENERGY-EFFICIENT SMART CARS WITH ADVANCED VOICE CONTROL

VIDEO TRANSCRIPT

Speaker 1: Tim Shea, a research scientist at Accenture Labs, has been working with an automotive client on prototyping the use of neuromorphic computing to help people interact with smart vehicles.

Speaker 2: Consumer demand for AI driven experiences is increasing rapidly, especially in the automotive industry. Customers expect responsive voice, gesture, and contextual intelligence from their vehicles. But current AI hardware is too power hungry, which can impact vehicle performance and limit the possible applications. Smart vehicles need more efficient edge AI devices to meet the demand.

Using edge AI devices to compliment cloud-based AI could also increase responsiveness and improve reliability when connectivity is poor. So we've built a proof of concept system with one of our major automotive partners to demonstrate that neuromorphic computing can make cars smarter without draining the batteries. We're using Intel's Kapoho Bay to recognize voice commands that an owner would give to their vehicle. The Kapoho Bay is a portable and extremely efficient neuromorphic research device for AI at the edge.

We're comparing that proof of concept system against a standard approach using a GPU. To build the system, we trained spiking neural networks to differentiate between command phrases. Then we ran the trains networks on the Kapoho Bay. We connected the Kapoho Bay to a microphone and a controller similar to the electronic control units that operate various functions of a smart vehicle.

We're targeting commands that reflect features that can be accessed from outside a smart vehicle, such as park here, or unlock passenger door. These functions also need to be energy efficient, so the vehicle can remain responsive even when parked for long stretches of time.

As a first step, we trained the system to recognize simple commands, such as lights on and lights off, open door, close door, or start engine. Using a combination of open source voice recordings and a smaller sample of specific commands, we can approximate the kinds of voice processing needed for smart vehicles. We tested this approach by comparing our train spiking neural networks running on Intel's neuromorphic research cloud against a convolutional neural network, running on a GPU.

Both systems achieved acceptable accuracy recognizing our voice commands, but we found that the neuromorphic system was up to a thousand times more efficient than the standard AI system with a GPU. This is extremely impressive and it's consistent with the results from other labs, as Intel will show further in their session on benchmarking the Intel OAE.

The neuromorphic system also responded up to 200 milliseconds faster than the GPU. This dramatic improvement in energy efficiency for our task comes from the fact that computation in Loihi is extremely sparse. While the GPU performs billions of computations per second, every second, the neuromorphic chip only processes changes in the audio signal and neuron cores inside low Loihi communicate efficiently with spikes.

This project demonstrates that neuromorphic systems can prove more efficient and more responsive than conventional solutions for AI in smart vehicles. This research is helping our partners in the automotive industry understand how Intel's neuromorphic systems might impact their next-generation products. And it helps us develop a roadmap for future neuromorphic applications.