Paving a New Path to AI-Driven Neuroscience

Researchers from McGill University and the University of Montreal are decoding the brain using neural networks and the Dell EMC Zenith supercomputer.

Scientific Research | United States

Business needs

Researchers need leading-edge artificial intelligence and advanced computing resources to accelerate brain mapping with a goal of understanding, and help pave the way for Alzheimer's and dementia treatments.

Solutions at a glance

- Dell EMC™ PowerEdge™ servers
- Intel® Xeon® processors and Optane® memory
- Intel® Omni-Path Architecture
- Dell EMC Ready Solutions for HPC NFS Storage
- Dell EMC Ready Solutions for HPC Lustre Storage
- Dell EMC Isilon F800 all-flash NAS storage

Business results

- Powering groundbreaking brain research
- Opening new frontiers in artificial intelligence
- Enabling work with enormous datasets
- Accelerating the training of large AI models

For their initial benchmarks on a public dataset, the research team worked with

20TB of data

Researchers will use MRI machines to record the activity in the brains of each subject for

500 hours

Reading minds in a brain lab

In the development of artificial intelligence applications, the holy grail is the creation of an artificial neural network that functions like the human brain. This is an elusive goal, because the human brain is an extremely complex organ that functions in flexible and fluid ways that can be difficult to replicate in the world of artificial intelligence (AI).

Today, researchers from McGill University and the University of Montreal are making breakthroughs in this area by using functional magnetic resonance imaging (fMRI) of the brains of people carrying out various cognitive tasks. The goal is to better understand and create computational models of how the brain works, and then use those models to train artificial neural networks to map the images to actions quickly and accurately.

For example, having a fully developed computational model of how memory works would make it possible to compare brain activity with what works or does not work and understand which model is playing out in the simulated brain of a patient. With this base, the research team could gain deep insight into the mechanics of memory function in those suffering from age-related brain illnesses, including Alzheimer's disease and other forms of dementia.

This would be a big leap forward for the AI world, according to one of the lead researchers on the project, Dr. Pierre Bellec, an associate professor at the University of Montreal. Dr. Bellec is the scientific director of the Courtois Project on Neuronal Modelling (NeuroMod), which is spearheading the collaborative research effort.

"Something the brain does really well is to switch from one context to another," Dr. Bellec explains. "It has very elaborate organization, and specialized networks and subnetworks, and those networks and subnetworks are able to reconfigure dynamically. By contrast, current architectures used by AI researchers are extremely specialized for certain types of tasks, and have a hard time generalizing over different contexts."

The researchers hope that by mimicking the architecture of the human brain they can develop a more versatile AI model that can generalize over different tasks, much the way the human brain does.

"AI has been an inspiration forever, but here we are not just drawing general principles, we are doing extensive imaging to map out the activity of the brain in an unprecedented level of detail," Dr. Bellec explains. "We're hoping to be able to draw directly from rich data to gain insight on how the brain works, rather than drawing from general, vague principles."

To collect the datasets for this ambitious effort, the research team has recruited a small group of volunteers — a half dozen subjects — to watch videos, look at images and play video games while they are in an MRI machine. The research team had to build a new game controller without any metal, printed in 3D plastic with a fiber optic cable connection. The machine allows the researchers to track and record the activity in the brains of the subjects as they carry out their tasks. The research team expects to gather many terabytes of data over the course of the five-year study, during which time each subject will spend around 500 hours in an MRI machine.

The NeuroMod project

The Courtois Project on Neuronal Modelling, also known in AI and computational neuroscience circles as NeuroMod, is a privately funded research effort based at the Montreal Geriatric University Institute. Neuromod aims to build more robust artificial intelligence systems that can replicate patterns in the brain as they have been measured with MRI data from human brain activity when learning, playing video games, looking at images and more. The project is pushing several barriers, from developing new neuroimaging and processing techniques to developing new AI frameworks to support these. NeuroMod data will also serve as a public resource for further research in neuroscience and AI.

"Essentially, we are trying to find a new way to integrate activity from human neural networks to help train artificial networks," Dr. Bellec says. "The hope is that if we manage to do that, we can create computational models of how the brain works. And potentially we can train new artificial neural networks that may perform better in some settings than what we have now."

The computing infrastructure and initial benchmarks

To move this project forward, researchers from the University of Montreal teamed up with researchers from Alan Evans' lab at McGill University who have extensive experience in high performance computing and work with MRI images that require large memory capacities.

They also sought the help of Dell Technologies and Intel, along with the data science and supercomputing resources of the Dell Technologies HPC & AI Innovation Lab in Austin, Texas. The team is using the lab's Intel-based Zenith cluster, which includes hundreds of Dell EMC PowerEdge™ servers with Intel® Xeon® Scalable Processors and the Intel® Omni-Path Architecture.

"People have done a lot of training on GPUs," Dr. Bellec says. "We wanted to use CPUs for this project, thinking a CPU deep learning architecture with a large memory capacity would be better for larger files. We thought this type of hardware would be a perfect fit for our use case, so we decided to run some benchmarks."

After testing on a GPU architecture, the team found that a CPU-based model can maintain similar performance — with validation accuracy reaching 99 percent after 10 epochs in distinguishing five types of body movements, and 91 percent after 20 epochs in classifying eight types of visual working-memory tasks. And while achieving performance similar to the GPU-based models, the CPU-based model requires much less training time — 20 minutes vs. 3 hours per epoch — when using 10 CPU nodes and two GPU cards, respectively. Considering CPU resources can often be more easily accessed, the project provides a more feasible solution for the application of deep neural networks on large-scale neuroimaging data by training the model directly on CPU hubs instead of waiting for other resources.

For those initial benchmarks, which kicked off the NeuroMod project, the research team used publicly available datasets from the Human Connectome Project, according to the team's senior postdoc, and team lead on the project, Dr. Yu Zhang.

The Human Connectome Project, which is mapping the neural connections in the brain, offers researchers access to fMRI scans from 1,200 subjects.

"Currently, we're using that project data and trying to decode brain cognition," Dr. Yu Zhang says. "The neural networks take a short series of 3D fMRI volumes and try to predict the specific task the subject was performing during the scanning."

In the benchmarks on the public dataset, the team is evaluating the performance of two different architectures — a traditional convolutional neural network and a more complex one, called ResNet, that has been used a lot in image processing.

"The benchmarks we are doing are what we call brain decoding, or in layman's terms, mind reading," Dr. Bellec says. "You look at brain images and try to predict what people were doing. In the next stage, we will have our own data, in which we play video games. The idea is to try to train an artificial neural network to play the video game in the style of the particular player."

Big data and big memory

By its nature, NeuroMod is a project that needs an HPC system with a large memory capacity to handle terabytes of data. In the initial benchmarks on the public dataset from the Human Connectome Project, the research team was dealing with nine terabytes of compressed data, or 20 terabytes in an uncompressed form.

"And that's only for the brain images," Dr. Bellec says. "We are also collecting data from all the videos, and we are collecting a lot of physiological data on the subjects, including heartbeat, respiration and the small motion of their eyes in very high resolution. All those auxiliary data, which we will eventually use in the model, can add tens of terabytes of data."

This deluge of data makes it all the more important to have ready access to an HPC cluster with big memory, which is what the team is getting through the Dell Technologies HPC & AI Innovation Lab and its Zenith supercomputer. This all-CPU system is on the Top 500 list of the world's most powerful HPC machines, and it has been designed to support massively parallel traditional scientific applications as well as emerging machine learning workloads.

Central to the Zenith system's capabilities are a variety of Intel® Xeon® processors with both host and coprocessors sporting Intel® Xeon® Platinum and Gold processors. In addition, the system marks a first for Dell Technologies as part of the Intel

Ready for a deeper dive?

- See Functional Annotation of Human Cognitive States Using Deep Graph Convolution by Yu Zhang, Loic Tetrel, Bertrand Thirion and Pierre Bellec.
- Dell Technologies HPC & AI Innovation Lab slidecast "Brain Decoding: Using Neural Networks to Read Minds," with Data Scientist Dr. Luke Wilson.
- For a closer and more technical look at the NeuroMod project, see the white paper "Bridging the gap between perception and action: the case for neuroimaging, AI and video games" by Dr. Pierre Bellec and Julie A. Boyle.

Fabric Builders program, which pushed the envelope on workloads that span HPC and AI with the Intel® Omni-Path Architecture (Intel® OPA) host fabric.

"We got access to the cluster in June 2019, and it's been very productive for us," Dr. Bellec says. "We have been able to run a number of benchmarks that we had not been able to run prior to that time. And this is just the beginning, or so we hope. We're getting familiar with the computing hardware architecture for deep learning, which we plan to use for years to come."

This is just part of the process when a research team is breaking new ground. Oftentimes, organizations need new HPC architectures that are built for the challenges of huge datasets and unique workloads.

"Many people are excited about being able to evolve neural networks in ways that are inspired by biology, and it's increasingly clear that we need a different type of hardware to do that," Dr. Bellec says. "And that's what we have with the Zenith cluster in the Dell Technologies HPC & AI Innovation Lab."

Learn more about high performance computing

Unlock the value of data with artificial intelligence

Share this story