HIGH PERFORMANCE DATA ANALYTICS

Pacific Northwest
NATIONAL LABORATORY

Proudly Operated by Battelle Since 1965

Message Passing Interface-based Machine Learning

CHALLENGE

Machine learning and data mining (MLDM) algorithms are becoming the standard in model generation from large-scale data sets. As the data produced by handheld devices, and other sensors continue to increase, it is becoming important to consider algorithmic and architecture advancements for reducing time-to-solution. Scaling MLDM algorithms to distributed memory architectures, especially connected with high-speed interconnects and state-of-the-art compute devices, is equally essential to address the ever-increasing volume of data sets. Several attempts to design distributed memory algorithms have been proposed, such as Distributed TensorFlow, Caffe (with message passing interface [MPI]), and Spark. However, it is necessary to consider distributed memory optimizations in conjunction with the algorithm to minimize (and, in a few cases, completely overlap) the communication time. Most importantly, any distributed memory changes should be seamless to an application writer, who instead should be concerned about developing a novel MLDM algorithm and managing distributed memory implementations.

Combining algorithmic and systems techniques for scaling out deep learning algorithms on very-large-scale systems.

CURRENT PRACTICE

We primarily focus on deep learning algorithms, including deep neural networks, convolutional neural networks, and recently proposed recurrent neural networks. There are a few widely used approaches for scaling MLDM algorithms, including CaffeOnSpark, FireCaffe, and Distributed TensorFlow (using general remove procedure call [gRPC]). A limitation of these implementations is that they do

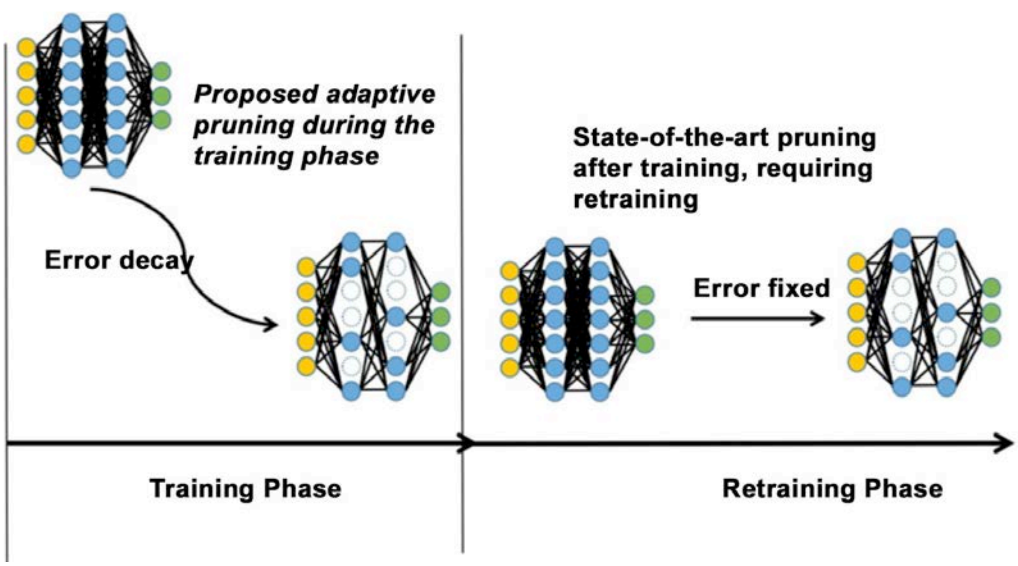

U.S. DEPARTMENT OF ENERGY

not leverage the high-performance computing (HPC) interconnect(s) natively. FireCaffe, which uses MPI, is not open source. Our intention is to consider a combination of algorithmic and systems approaches for scaling deep learning algorithms with little to no intervention from an analyst.

TECHNICAL APPROACH

Our technical approach consists of a combination of algorithmic and systems techniques for scaling out deep learning algorithms on very-large-scale systems. Specifically, we are considering the following design choices:

- **Scaling Caffe on central processing unit (CPU) or graphics processing unit (GPU) clusters using MPI.** While several approaches have been proposed in the literature, we will leverage the fault tolerance project to provide scalable and fault tolerant deep learning implementation, presently not provided by other high-performance implementations.
- **Addressing heterogeneous devices.** Currently, most implementations either leverage CPUs or GPUs. We are working on a design that may leverage them simultaneously.
- **Pruning redundant/noncontributing neurons.** Several approaches have been considered in the past for removing neurons/synapses after the network has been trained. However, these approaches are problematic because they increase the training time. We are considering approaches where the training time can be reduced, without deminishing the solution's accuracy.
- **Deep learning on deep memory hierarchies.** We are developing solutions to leverage the deep memory, including volatile/non-volatile memories, and optimizing deep learning algorithms for deep memory hierarchies, e.g., for file systems such as Lustre.
- **Adaptive runtimes for flexible deep learning computations.** We are working on runtimes that can automatically shrink or expand, depending on the nature of the computation. A few of these properties are generated by the adaptive nature of the algorithms featured as part of this project.

IMPACT

The potential impact of this research is significant in terms of adopting distributed memory architectures seamlessly. With adaptive runtimes, user would be able to reduce pressure on compute resources while scaling down the training time by removing unnecessary samples and neurons—all automatically.

Contacts

Abhinav Vishnu
Principal Investigator
(509) 372-4794
Abhinav.Vishnu@pnnl.gov

John R. Johnson
Program Director
(509) 375-2651
John.Johnson@pnnl.gov

Pacific Northwest
NATIONAL LABORATORY
Proudly Operated by **Battelle** *Since 1965*