US-ATLAS DISTRIBUTED LEARNING WORKING GROUP PROGRESS REPORT

March 21, 2019

Abid M. Malik, Meifeng Lin, Wei Xu

Introduction

The High Luminosity Large-Hadron Collider (HL-LHC) will deliver data from 10 times the luminosity compared to the LHC, bringing quantitatively and qualitatively new challenges due to event size, data volume, and complexity. The physics reach of the experiments will be limited by the physics performance of algorithms and computational resources. Machine learning (ML) applied to particle physics has the potential to provide improvements in both of these areas. However, producing fast and accurate machine learning models with large data sets require large computational resources.

Machine Learning Challenges

The working group is looking into the machine learning challenges for the HEP (High Energy Physics) community from the following two perspectives:

- Frameworks for distributed machine learning for data parallelisam.
- Readiness of PanDA Framework to take HEP distributed machine learning loads.

Distributed Machine Learning

The HEP community is considering $Horovod$ and $MPI-learn$ frameworks for distributed learning from LHC. Horovod, developed by Uber, is a distributed framework for Tensorflow, PyTorch, and Keras. MPI-Learn, developed and supported by CERN-Openlab, is a lightweight MPI-Based Python framework for distributed training for Tensorflow, PyTorch and Keras. We are doing a detailed performance analysis for both frameworks using 3D GANs. We are using Score-P Performance Analysis Tool to analyze the computation and communication characteristics of both frameworks. We are also doing manual instrumentation of the Python code to analyze per GPU/CPU node performance. We did some initial runs using Institutional Cluster(IC) at BNL with 20 $P100$ GPUs.

Current Status

We got access to the Summit machine in Feb. 2019. The Summit software environment does not support Tensorflow and Horovod yet. We prepared the software stack on Summit for Tensforflow 1.12 and Horovod. Most of these packages had to be compiled from source due to the unique software environment on Summit. With the software frameworks we installed, we ran successfully machine learning models such as MNIST and ResNET on Summit using 200 V100 GPUs.

Currently, we are porting 3D GANs on Summit. For experimentation and analysis, we are closely working with Amir Farbin (ATLAS ML Coordinator) from the University of Texas at Arlington, and Dr. Jean-Roch from Caltech. Scientists from other National Labs (LBNL and ANL) are also participating in the discussion during the bi-weekly meetings of the group. The outcome of this work will help explain the performance bottlenecks of the two frameworks under the HEP distributed machine learning workloads.

Readiness of PanDA for Machine Learning Loads

To get the best inference performance from a machine learning model, its design parameter need to be tuned. This is done by hyper-parameter optimization. Hyper-parameter optimization requires large data (several millions of inputs) and training of a lot of models. This mean, a lot of jobs (10^5) to be submitted for each model optimisation. Each job can take as much as 2 days on a single node or multiple nodes. The PanDA workload manager which has been proven to scale very efficiently on the Leadership Computing Facilities can be used for hyper-parameter optimization problem for single as well as distributed machine learning models. The PanDA workload manager can be used to monitor and collect run-time information of training models. The information can be used by an $Optimizer$ to pick the the best model off-line or improve the run-time search strategy on the fly. As a first step, the group decided to do a performance analysis of the framework with respect to the machine learning loads using Horovod and MPI-Learn. The group did some initial experiments using IC.

Current Status

Currently, the group is planning to do more detailed analysis using Summit. The group is working with Kaushik De from the University of Arlington, and Alexi Klimentrov, Torree Wenaus, and Sergey Panitkin from BNL. Scientists from other National Labs, who are part of the PanDA group, also participate in the discussion during the bi-weekly meeting.

Future Direction

Hyper-Parameter Optimization

Current supercomputers are powerful systems, but present challenges when faced with problems requiring large machine learning workflows. The PanDA framework has the potential to provide the same services to the HEP community for its machine learning loads which the CANDLE framework is providing to the cancer research for machine learning loads. PanDA can be a Scalable Workflow Framework for Machine Learning Loads for HEP.

Other HEP groups, such as DUNE, have also shown keen interest in the ML workloads.

Provenance

Can data provenance (or a similar concept) help explain why a deep learning model produces a particular output for a given input? Since deep learning models are thought to be black-box functions, It would be interesting to see whether provenance information could give insights on how deep learning models work. Reproducibility of the results will be a big issue for the simulation codes using machine learning models.

Scalability with Large Batch Size

Scaling distributed machine learning with large number of nodes with a good rate of convergence is still a big research challenge for the HPC community working in this area. Research in new communication All-REDUCE algorithms and data distribution approaches to build a large batch size is needed .

Explainable Machine Learning

As machine learning models become deeper and wider, it is essential to understand how the networks make decisions, to control over their internal processes and impose domain knowledge. To solve this challenge, interactive visualization has shown promising results to enhance the interpretability of complex machine learning networks. A wide range of research has been dedicated to different aspects and stages of the models including model understanding, debugging and refinement [1]. As for GAN networks, GAN Lab [2] presents an interactively training environment for generative models and visualizes the dynamic training process intermediate results. GAN Dissection [3] is a way to analytically inspect the internal representations at the unit-, object-, and scene-level of a generative adversarial network to understand how internal units align with human-interpretable concepts. DGMTracker [4] is a visual analytics tool that helps experts understand and diagnose the training processes of deep generative models. It employs the blue noise sampling algorithm and credit assignment algorithm to detect which portions of the input images cause a training failure for a particular image set. Therefore, research works are required to leverage these existing methods and extend to distributed computing environment on scientific problems.

Bibliography

[1] J. Choo and S. Liu. Visual analytics for explainable deep learning. *IEEE Computer Graphics and Applications*, 38(4):84–92, Jul./Aug. 2018.

[2] Minsuk Kahng, Nikhil Thorat, Duen Horng Chau, Fernanda B. ViÃl'gas, and Martin Wattenberg. Gan lab: Understanding complex deep generative models using interactive visual experimentation. *CoRR*, abs/1809.01587, 2018.

[3] David Bau, Jun-Yan Zhu, Hendrik Strobelt, Zhou Bolei, Joshua B. Tenenbaum, William T. Freeman, and Antonio Torralba. Gan dissection: Visualizing and understanding generative adversarial networks. *arXiv preprint arXiv:1811.10597*, 2018.

[4] M. Liu, J. Shi, K. Cao, J. Zhu, and S. Liu. Analyzing the training processes of deep generative models. *IEEE Transactions on Visualization Computer Graphics*, 24(1):77–87, Jan. 2018.