Hyper-parameter Tuning on Titan

Achievement: CDA's NIML team achieved a peak of 20 petaflops on Titan (single precision) using MENNDL and RAvENNA, two different software codes for optimizing deep learning hyper-parameters.

Significance and Impact: MENNDL and RAvENNA are two codes developed by members of CDA's Nature Inspired Machine Learning team to address the challenge of tuning a deep learning convolutional network to a data set of interest. Using real world data, these codes achieved a peak of 20 petaflops on Titan, using single precision. Currently, no other known deep learning code has achieved this. This allows time to solution (tuning hyper-parameters) to be reduced from months to hours.

Research Details:

- Fully automated hyper-parameter tuning using MENNDL / RAvENNA allows deep learning to be applied to new, untouched, data sets with little to no domain knowledge, and taking just a few hours to perform rather than months of a domain expert manually tuning the parameters.

Sponsor/Facility: Work was performed ORNL and the ORNL OLCF. Work sponsored by DOE.

PI and affiliation: Robert Patton from CSMD, ORNL

Team: Steven Young, Travis Johnston, Thomas Potok

Publication: N/A

Overview: Deep learning has been applied to a wide range of real applications. For each application, a specific topology and parameter values are needed to define a deep learning network that can best classify the data set. Identifying this topology and parameter values is challenging and computational difficult. MENNDL and RAvENNA are two codes developed by members of CDA's Nature Inspired Machine Learning team to address this challenge. Using real world data, these codes achieved a peak of 20 petaflops on Titan, using single precision.