# Exploring Deep Learning for Science Benchmarks on DOE Supercomputers

Aristeidis Tsaris[1], Jacob Balma[2], Murali Emani[3], Steve Farrell[4], Geoffrey C Fox[6], Yin Junqi[1], Thorsten Kurth[7], Abid Malik[5], Prabhat[4], Mallikarjun Shankar[1], Venkatram Vishwanath[3]
[1]Oak Ridge National Lab, [2]Cray Inc., [3]Argonne National Laboratory, [4]Lawrence Berkeley National Laboratory, [5]Brookhaven National Laboratory, [6]Indiana University Bloomington, NVDIA[7]

## Introduction

**High-performance computing** is seeing an upsurge in workloads with data in unprecedented scale that require **data analysis**:

- Deep Learning models are used in several domains such as cosmology, particle physics and biology
- Problems typically include image detection, segmentation, synthetic data generation

It is critical to understand the **performance of deep learning models on HPC** systems on the exascale era and beyond.

**We are designing an HPC ML benchmark suite,** that will help:

- Better understand the model-system interactions
- Better understand the deep learning workloads, optimize them and identify potential bottlenecks
- Quantify the scalability for different deep learning methods, frameworks and metrics on hardware diverse HPC systems

## What's done

- Work is done within MLPerf-HPC working group
- Identify initial target systems: Summit@OLCF, Cori@NERSC, Theta@ALCF
- Select a list of candidate workloads. Not shown here, but in the works:
  - Language modeling on OpenWebText dataset: utilize communication patterns unique to sequence model-parallelism
  - Large scale multi-label image classification on TenCent dataset

## What's next

- Systematic study of the two benchmarks on the 3 target systems
- Finalize metrics collection and software tool for system performance (flops, I/O, communication) across all 3 systems
- Get the language modeling and the Image classification benchmark up and running
- Identify and pursue commonalities with main track MLPerf, best practices for data, model and pipelining parallelism
- Use our benchmarks to evaluate next-generation HPC systems
- Update suite with new benchmarks

## Climate segmentation

Segmentation of extreme weather phenomena in climate simulation (768 x 1152 x 16 image shape) **arXiv:1808.04728**

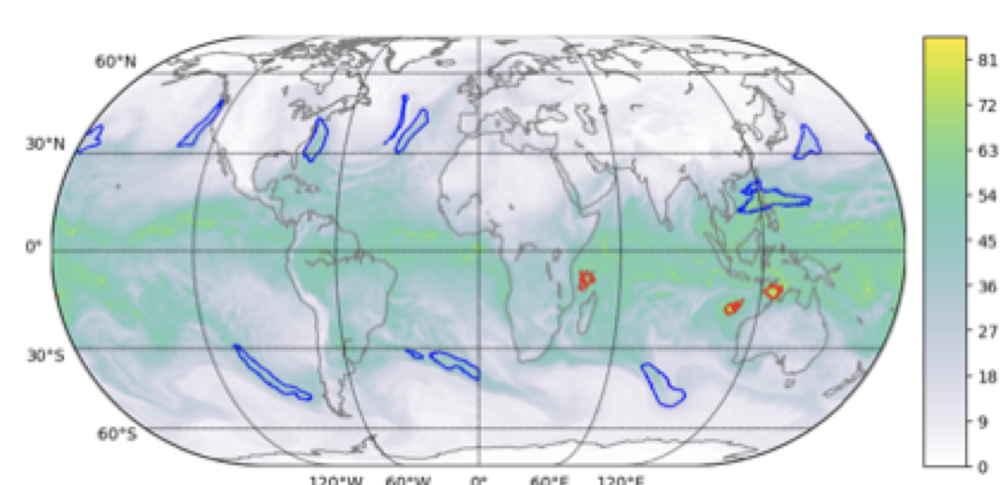

- High-dimensional inputs and multi-label targets
- 16-channel 2D images of size (768, 1152, 16), 3-class per-pixel
- **Large dataset (20 TB) stresses I/O subsystem**

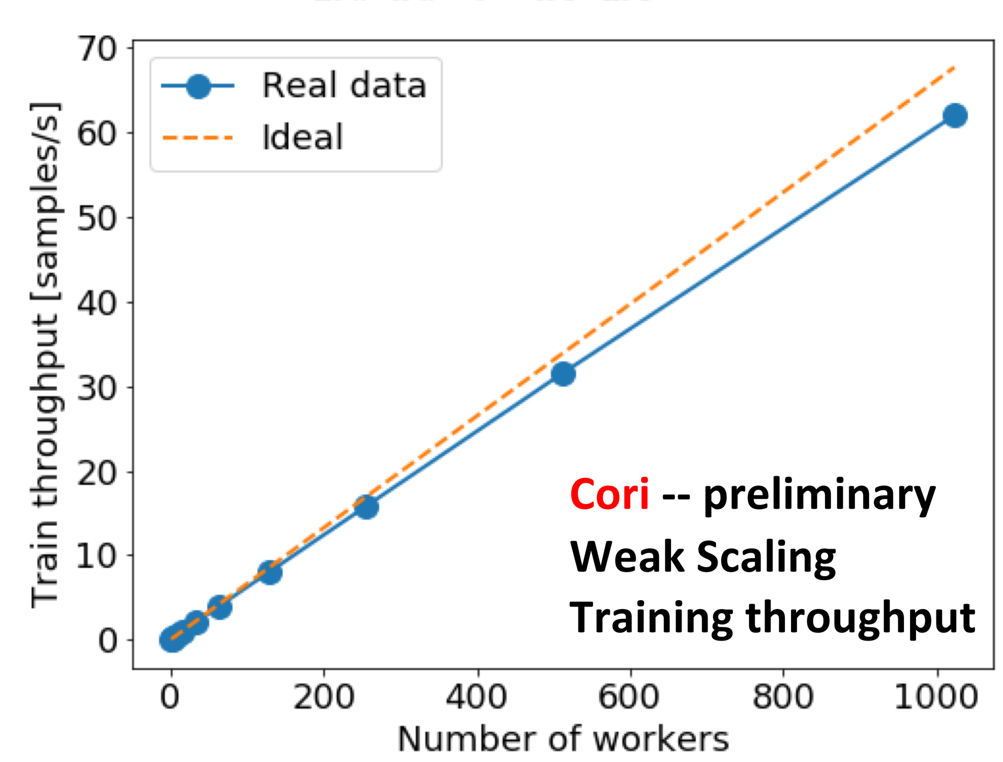


**Measurements on target systems with mini-dataset**

- **Cori:** 1024 KNL nodes, 1.33 PF/s peak, 1.14 PF/s sust **[FP32]**
- **Theta:** 1024 KNL nodes, 1.11 PF/s peak, 1.09 PF/s sust **[FP32]**
- **Summit:** 170 nodes (x6 v100s) 48.08 PF/s peak, 38.71 PF/s sust **[FP16]**
- Further analysis underway

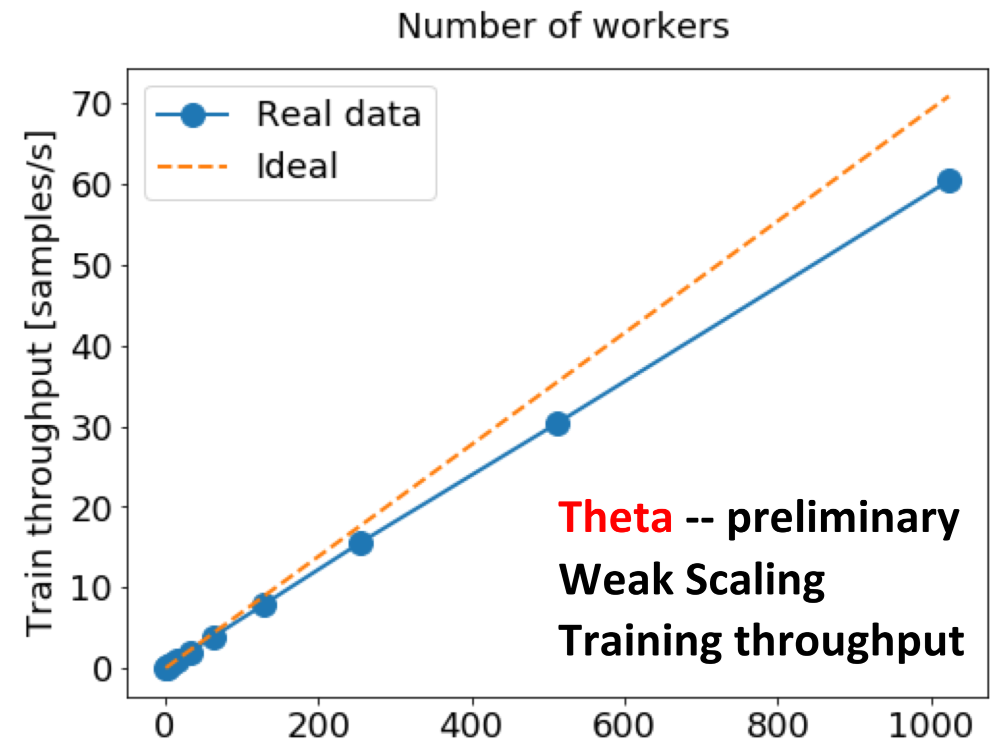


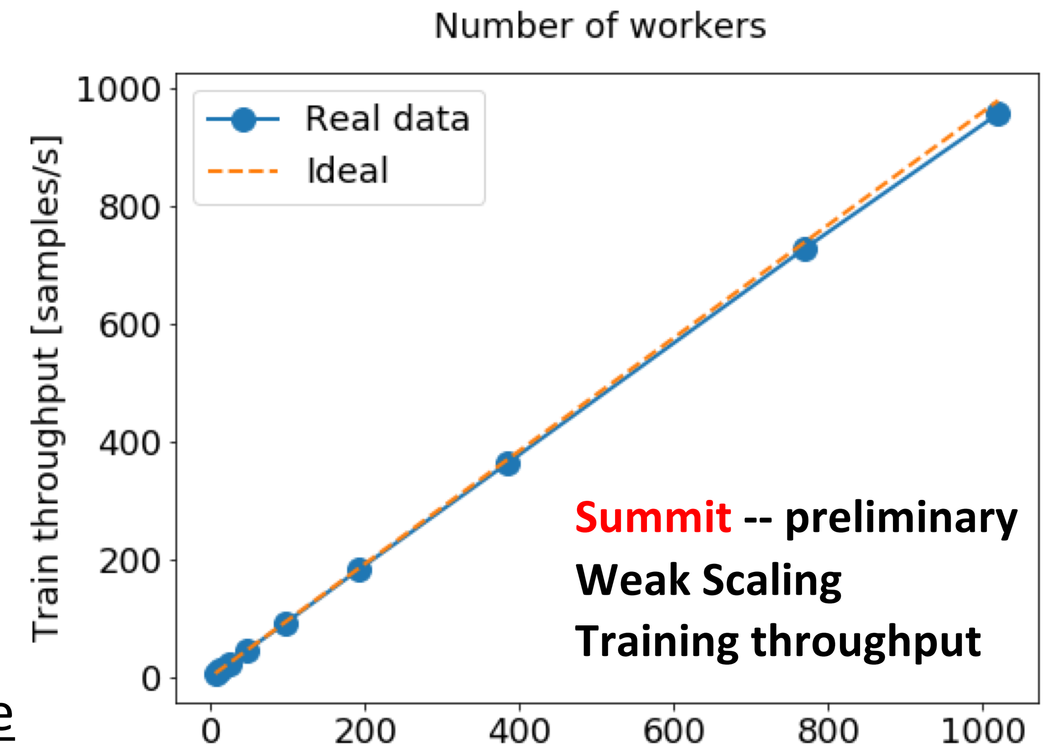


**Reference implementation:**
https://github.com/sparticlesteve/climate-seg-benchmark

## CosmoFlow

Predict cosmology parameters from whole-universe dark matter simulations (128 x 128 x 128 x 4 crops) **arXiv:1810.01993**

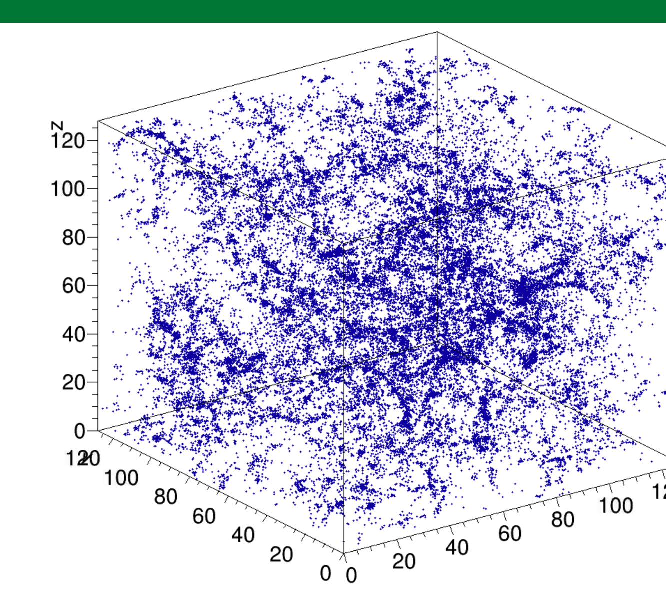

- Data-parallel implementation across 3D Convolutions
- **Multi-channel Volumetric input data stresses on-node bandwidths and memory capacities**
- **Communication patterns unique to convolutional model-parallelism**

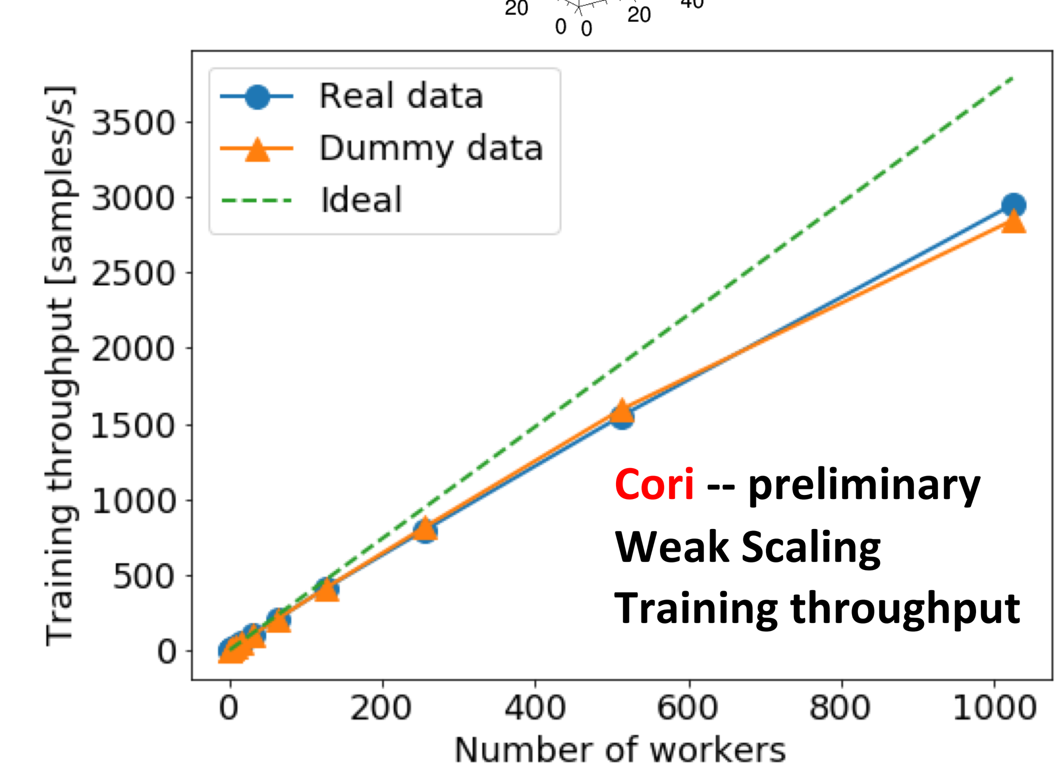


**Measurements on target systems**

- **Cori:** 1024 KNL nodes, --78% scaling eff
- **Theta:** 1024 KNL nodes, -- 84% scaling eff
- **Summit:** 170 nodes (x6 v100s), -- 72% scaling eff
- Further studies underway

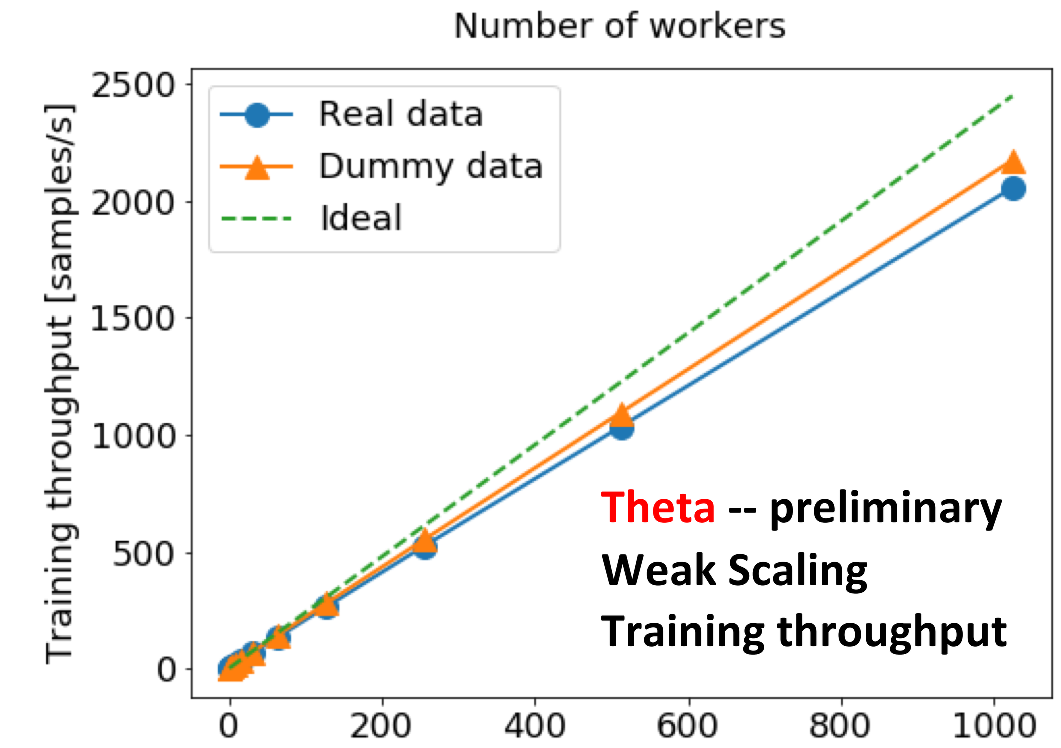


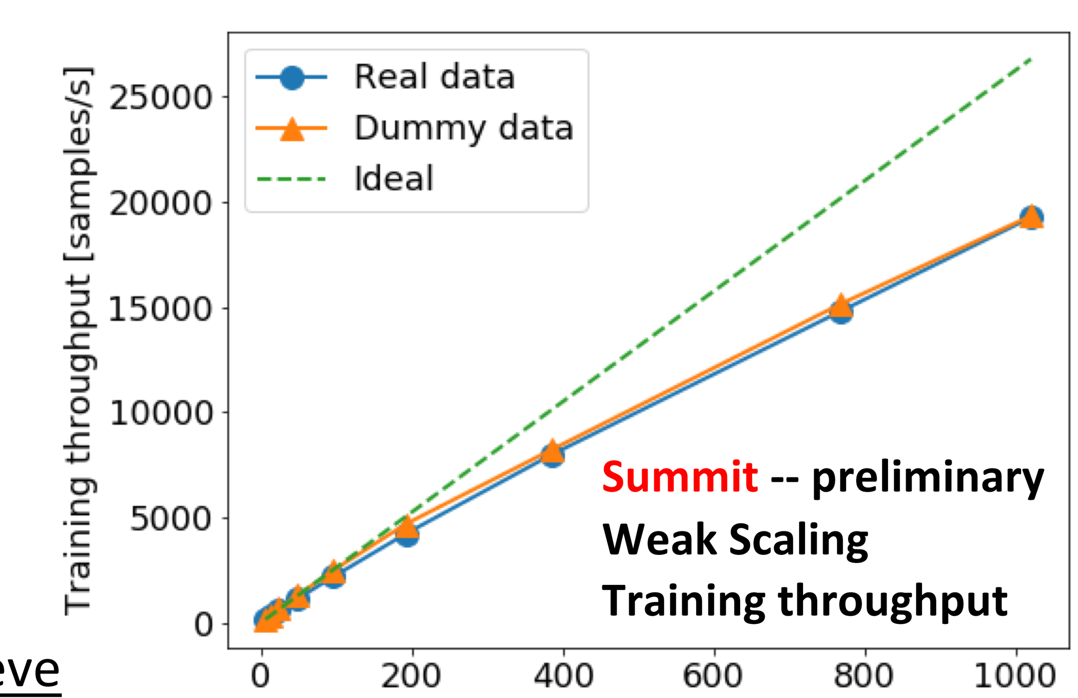


**Reference implementation:**
https://github.com/sparticlesteve/cosmoflow-benchmark