Scientific AI at Scale on the Perlmutter supercomputer at NERSC

Cray User Group Meeting
May 5th, 2022

Peter Harrington
Data & Analytics Services Group
NERSC

Motivation & Outline

- AI for science is maturing and can be transformative
- Cutting-edge requires supercomputing scale
- Enabling scientific AI at scale requires attention in:
    - Deploying systems and software
    - Developing applications
    - Empowering the community

HPC systems can enable cutting-edge AI applications

- Recent AI wave in DoE Science moving beyond proof of concept to maturity
  - *Transformative* performance will need re-framing of problems; large datasets and bigger models

- Role for HPC centers like NERSC
  - Productionized training, optimization, and error analysis on large models
  - Interaction with existing simulation codes and data pipelines

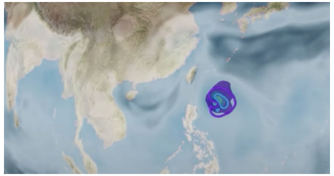

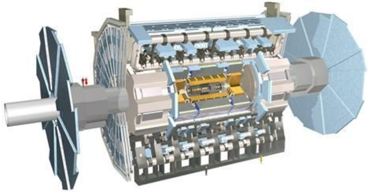

HPC systems can enable cutting-edge AI applications

- Work still needed
  - E.g. model development; methods for scaling; software tooling
  - Fast moving AI field where trends can strongly impact computational needs

- Develop through strong connections between HPC and AI applications
  - Application readiness program; in-depth engagements w/ scientists; user support & trainings

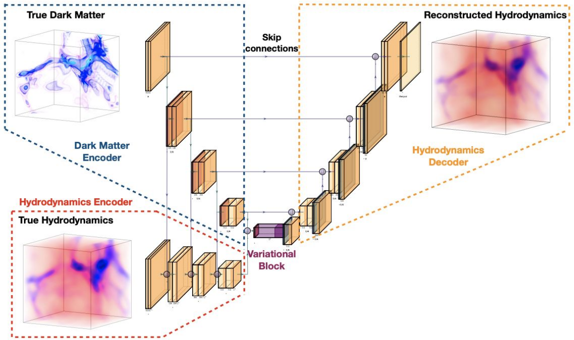

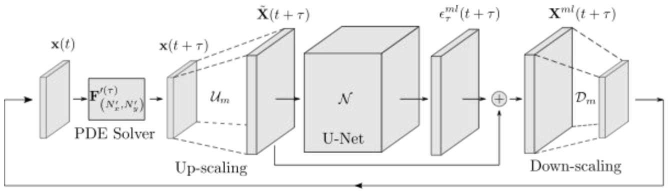

NERSC: Mission HPC for DOE Office of Science Research

Large compute and data systems

- Cori: ~12k CPU Nodes
- Perlmutter: ~6k A100 GPUs
- 128PB Community Filesystem ....

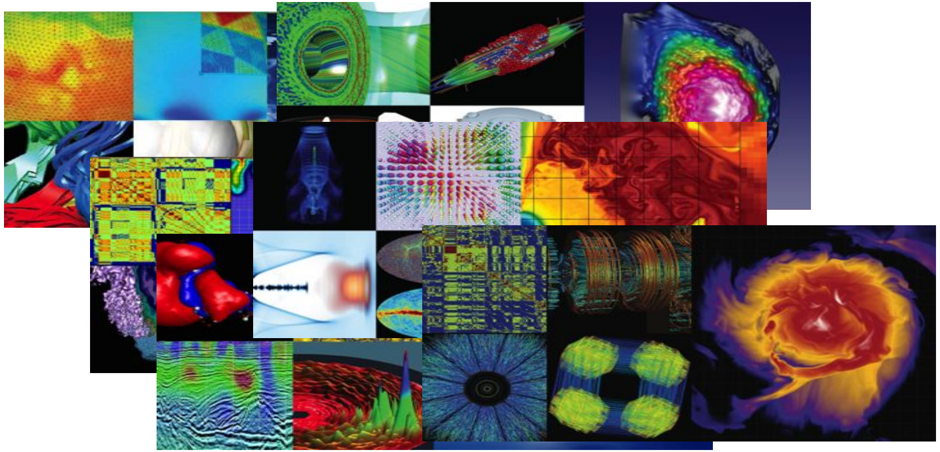

Broad science user base

- 7,000 users,
- 800 projects,
- 700 codes

https://www.nersc.gov/users/accounts/

NERSC AI Strategy

Deployment

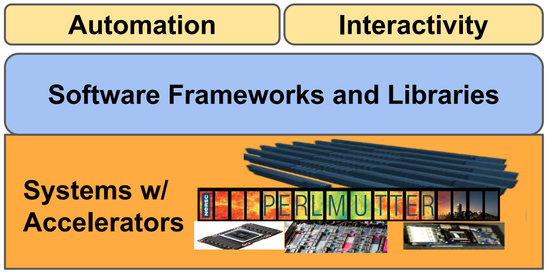

Methods and Applications

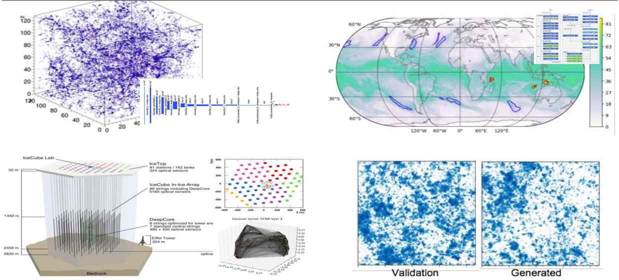

Empowerment

- ***Deploy*** optimized hardware and software systems
- ***Apply*** AI for science using cutting-edge methods
- ***Empower*** through seminars, workshops, training and schools

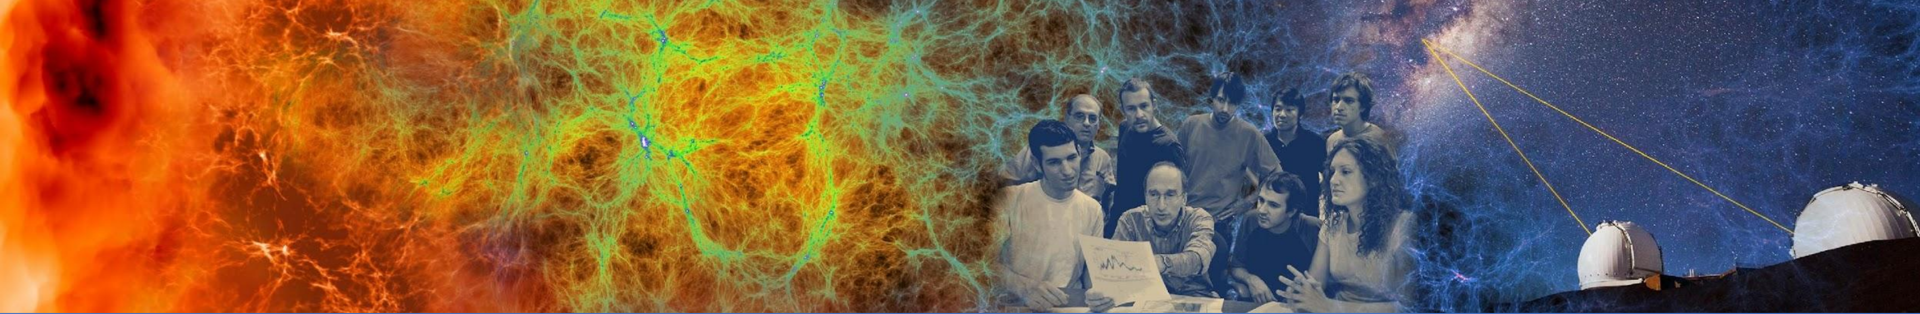

Deployment:
NERSC AI Systems, Software and Workloads

NERSC

BERKELEY LAB

U.S. DEPARTMENT OF ENERGY | Office of Science

Perlmutter: A Scientific AI Supercomputer

HPE/Cray Shasta system

Phase 1 (in early science phase):

- 12 GPU cabinets with 4x NVIDIA Ampere GPU nodes; Total >6000 GPUs
- 35 PB of All-Flash storage

Phase 2 (2022):

- 12 AMD CPU-only cabinets
- HPE/Cray Slingshot high performance network

Optimized software stack for AI

Application readiness program (NESAP)

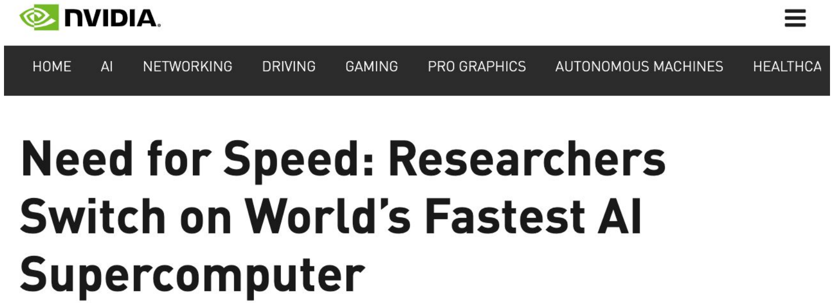
NVIDIA

HOME AI NETWORKING DRIVING GAMING PRO GRAPHICS AUTONOMOUS MACHINES HEALTHCA

Need for Speed: Researchers Switch on World's Fastest AI Supercomputer

See a growing scientific AI workload at NERSC

- We instrument user [python imports]
  - Users of DL frameworks increased more than 6x from 2018 to 2021
- Track ML trends through 2-yearly survey
  - Targets scientific communities which (potentially) use HPC resources (NERSC and non-NERSC users)
  - Covers problem type, workload, model architectures, framework, scaling, hardware, software, and usage of NERSC ML stack

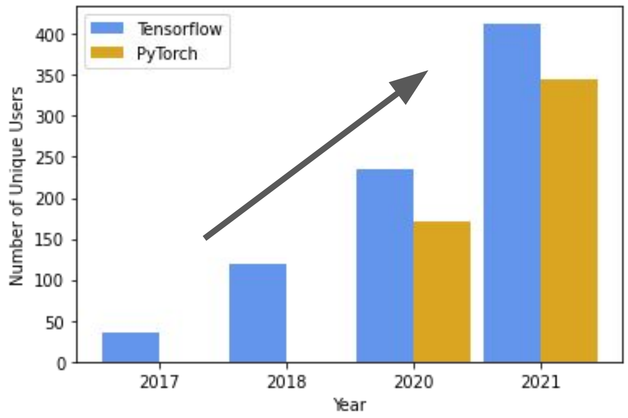

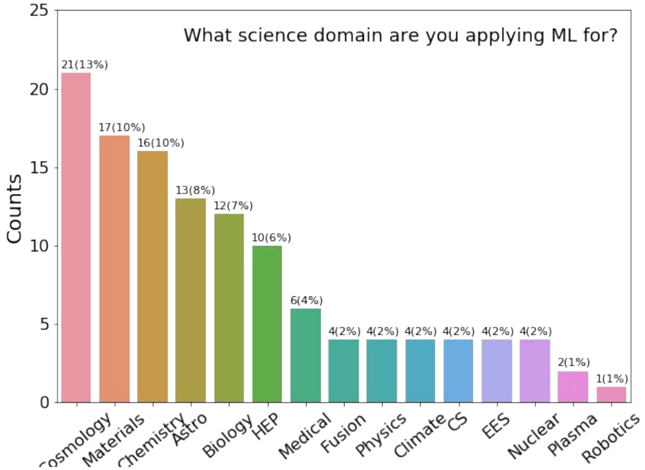

BERKELEY LAB

Software for scientific AI

- Observe typical spread over major AI/ML frameworks; trends for 2022 will be interesting

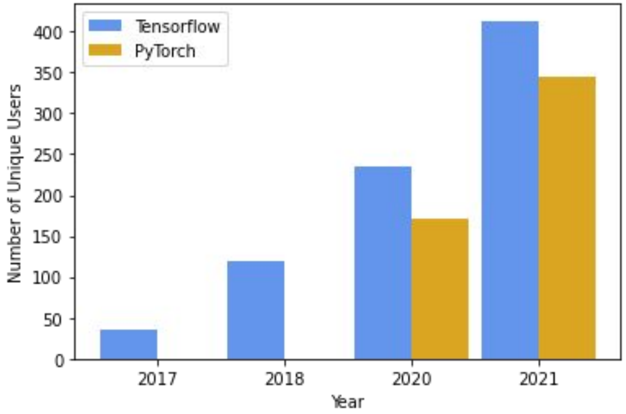

TensorFlow

Keras

PyTorch

Google Search trends:

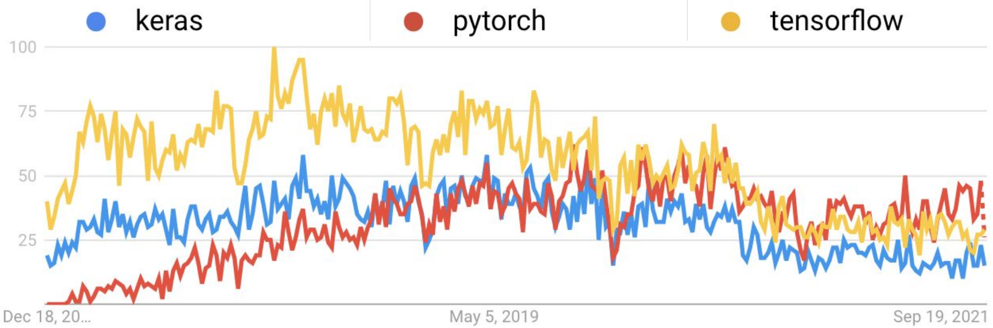

Need for resources and scale for scientific AI

- Scale allows rapid prototyping/model evaluation
- Volume of scientific datasets can be large and complex
- Data parallelism is currently the most common strategy in practice
- Horovod is the leading non-native parallelism framework.

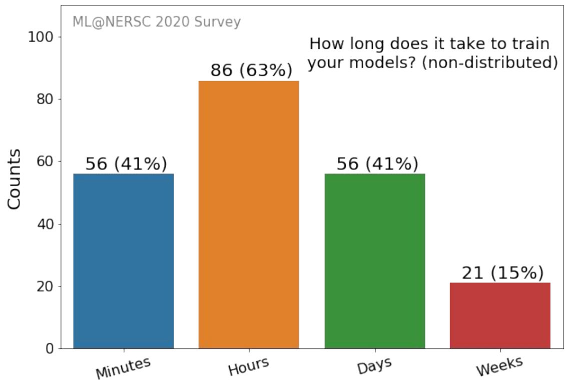

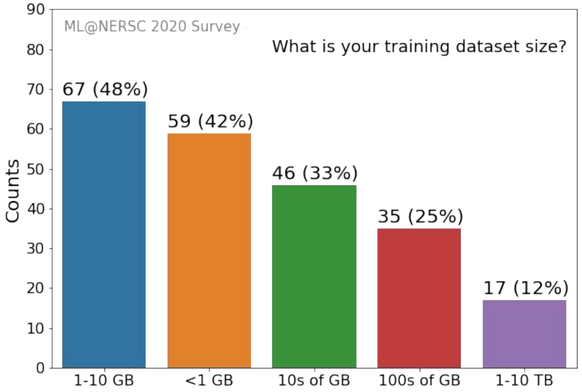

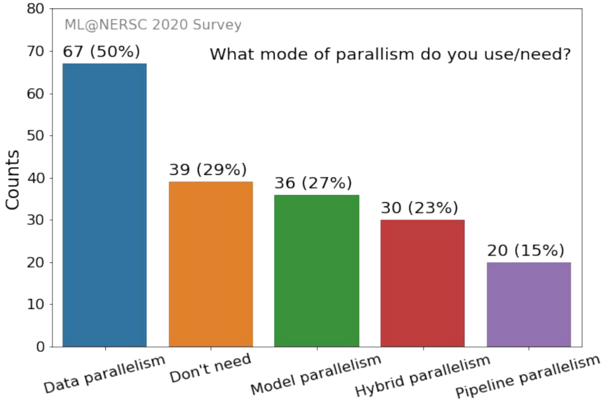

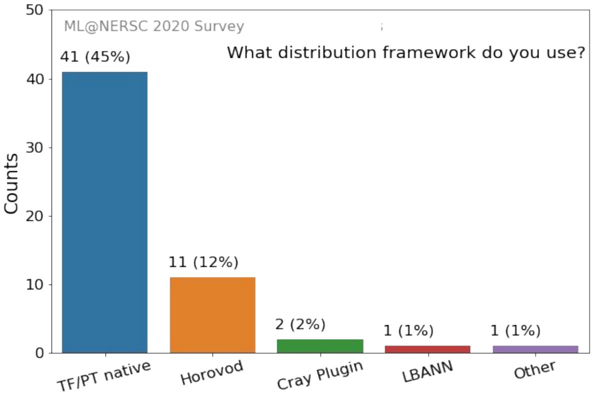

BERKELEY LAB

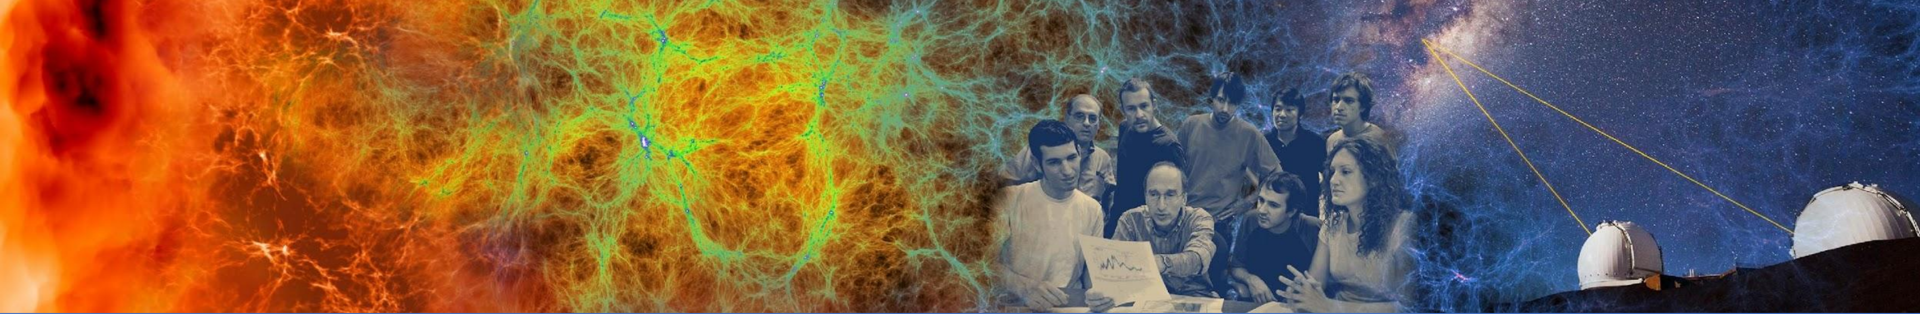

Lessons learned from NERSC deployments

NERSC

BERKELEY LAB

U.S. DEPARTMENT OF ENERGY | Office of Science

Scientists need performant and flexible software installations

Demand for:

- **Performant installations** of the most popular frameworks and libraries
- **Flexibility** for users to customize their solutions

On Perlmutter we chose to deploy both:

- Custom-built modules for TensorFlow, PyTorch
- NVIDIA's NGC containers
  - Container environment optimized for A100s and was crucial during deployment
  - Effectively debugged several minor deployment issues through close engagement with NVIDIA

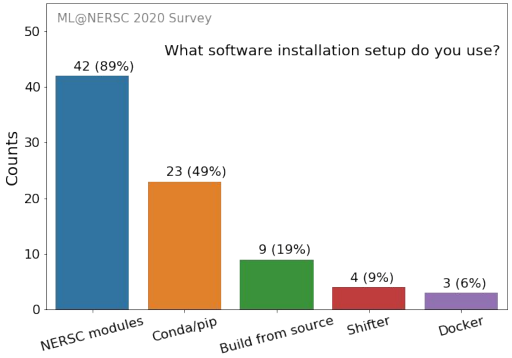

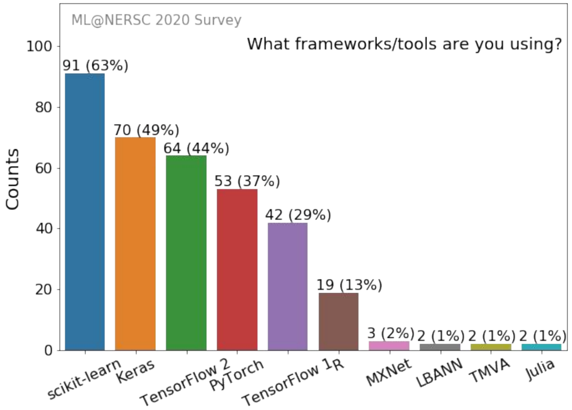

https://docs.nersc.gov/machinelearning/

Scientists need productive interfaces for experimentation

JupyterHub service provides a rich, interactive notebook ecosystem on Perlmutter

- Now over 2000 users at NERSC!
- A favorite way for users to develop ML code

Users can run their deep learning workloads

- on dedicated Perlmutter GPU nodes
- using our pre-installed DL software kernels
- or using their own custom kernels

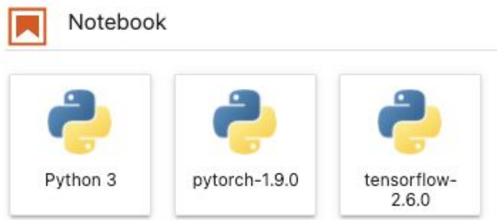

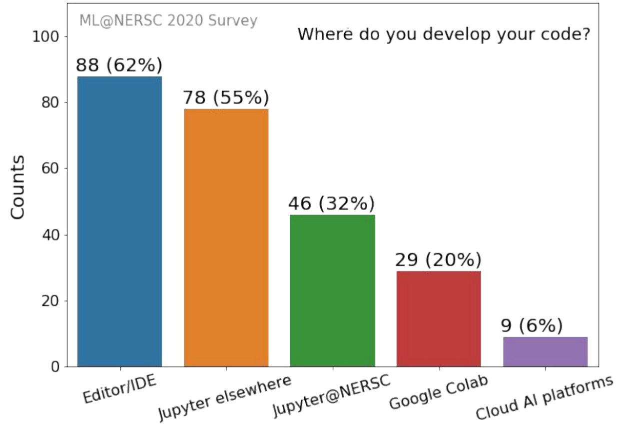

Scientific DL also needs HPC-enabled optimization tools

- Model selection/tuning is still critical for getting the most out of deep learning
- Computationally expensive: need for HPC
- Many methods and libraries exist for tuning model hyper-parameters
  - Enable users to use whatever tools work best for them
- Tools can need [adaption to work well on HPC]

RAY tune Weights & Biases

SIGOPT OPTUNA

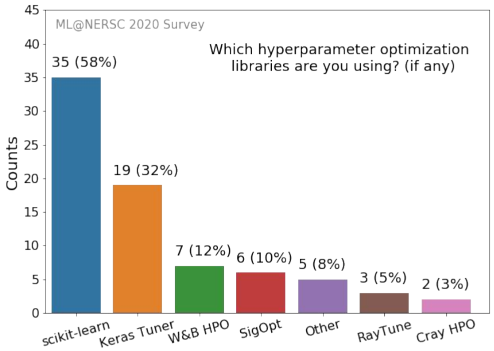

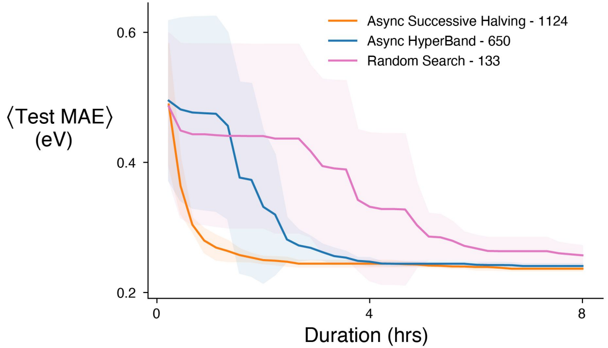

Multi-node RayTune HPO on Graph Neural Network models for catalysis applications ([B. Wood et al.])

AI compute performance requires benchmarking and tuning

MLPerf™ is the industry standard benchmark for ML performance

NERSC played active role to develop MLPerf HPC benchmark suite

- Built with *scientific applications* that push on HPC systems in important ways. Currently including:
  - CosmoFlow - 3D CNN predicting cosmological parameters
  - DeepCAM - segmentation of phenomena in climate sims
  - OpenCatalyst - GNN modeling atomic catalyst systems
- **MLPerf HPC v1.0 release at SC21 conference:**
  - Time-to-train and "Weak-scaling" throughput metrics
  - 31 submissions from 8 submitters on 9 diverse HPC systems
  - Best benchmark results improved by 4-7x from v0.7 round
  - Strong-scaling submission scale up to 2,048 GPUs
  - "Weak-scaling" submission up to 5,120 GPUs (Perlmutter) and 82,944 CPUs (Fugaku)
- Deeper analysis paper on v0.7 round at the SC21 MLHPC workshop

ML Commons

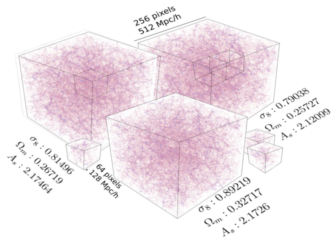

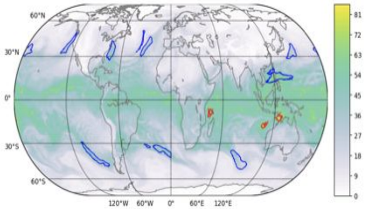

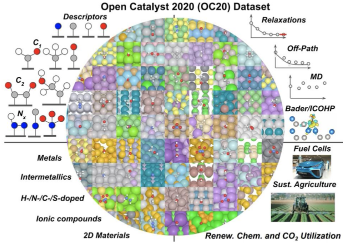

MLPerf HPC v1.0 - NERSC highlights and analysis

- Ran early in Perlmutter deployment
  - Very valuable to shake out system
- Leading time-to-train result for OpenCatalyst, second place results for CosmoFlow and DeepCAM
- Largest scale GPU throughput measurement (5120 GPUs)
- Subsequently performed in-depth profiling comparison with Selene
  - Dominant bottleneck from network (allreduce at 64 nodes ~3x slower)
    - Phase 2 of Perlmutter bringing slingshot 11 network upgrade!
  - Smaller 1 node difference from some unoptimized kernels and memory bandwidth on 40 vs 80G A100s

Perlmutter time-to-train results

Submitter	System	GPUs	Benchmark results (minutes)		
			CosmoFlow	DeepCAM	OpenCatalyst
LBNL	Perlmutter Phase 1	512			111.86
LBNL	Perlmutter Phase 1	1024	8.5		
LBNL	Perlmutter Phase 1	2048		2.51	

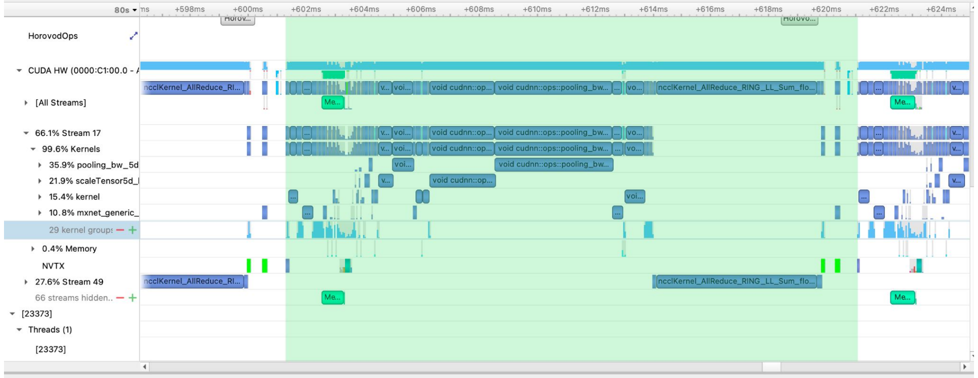

BERKELEY LAB

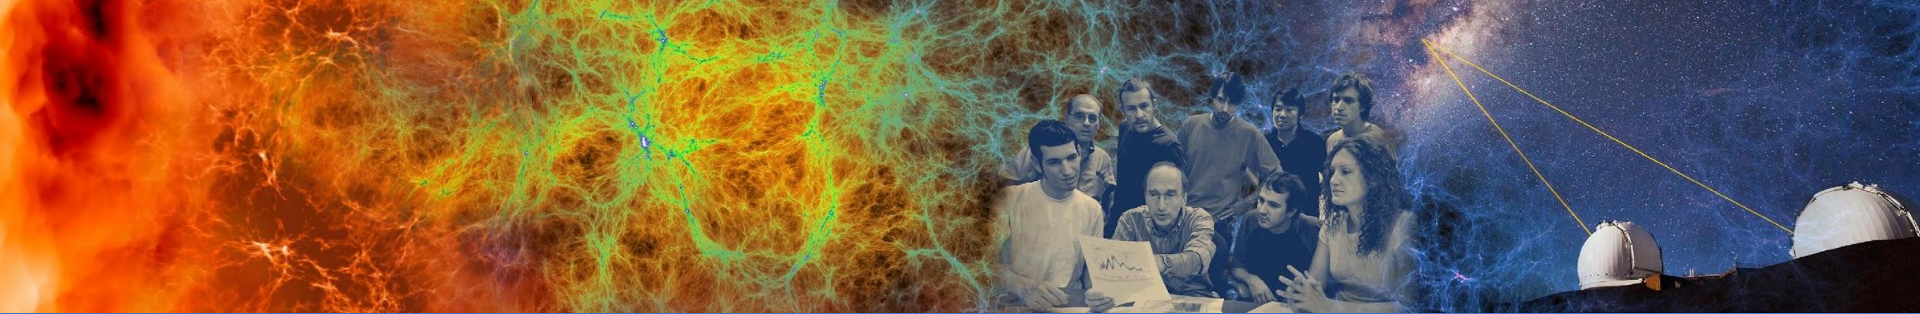

Applications and Empowerment: Powered By Perlmutter

NERSC

BERKELEY LAB

U.S. DEPARTMENT OF ENERGY | Office of Science

Accelerating science with AI

Extract

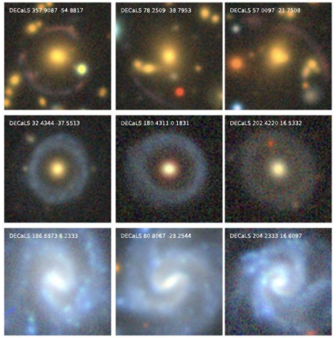

Hayat et al. 2021 arXiv:2012.13083

Enhance

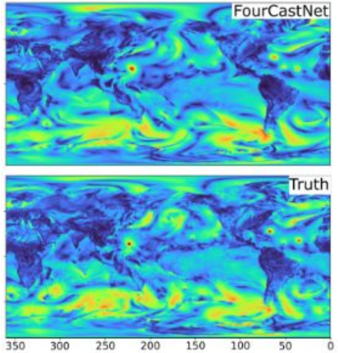

Pathak et al. 2022 arXiv:2202.11214

Explore

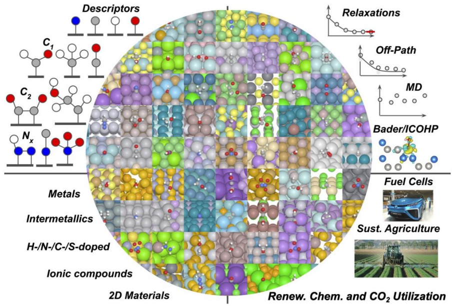

Chanussot et al. 2021 arXiv:2010.09990

Parallels with industry applications but current approaches increasingly incorporate science-specific structures

BERKELEY LAB

Extract: Self-supervised sky surveys

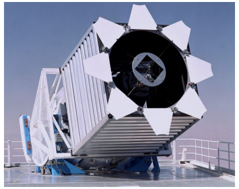

- Sky surveys image billions of galaxies that need to be understood
- Limited “labels”, so can learn in *semi-supervised* way
- Pre-training on entire dataset on HPC, downstream task can be on laptop/edge
- Recently used to find > 1000 previously undiscovered strong-lens candidates

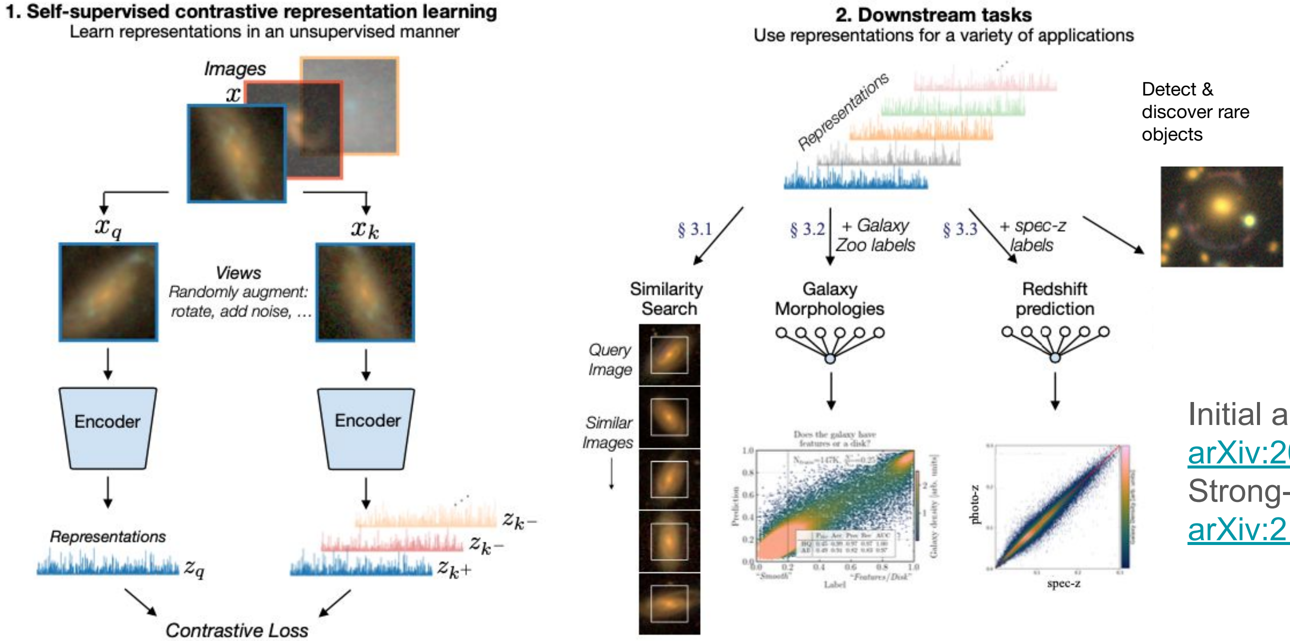

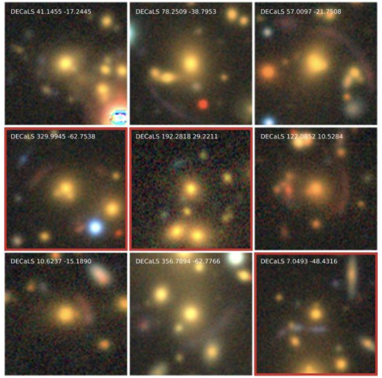

Initial approach: Hayat et. al. (2020)
arXiv:2012.13083
Strong-lens analysis: Stein et. al. (2021)
arXiv:2110.00023

Enhance: Data-driven atmospheric modeling

Pathak et al. 2022
arXiv:2202.11214

- Data-driven modeling of atmospheric flows using a state-of-the-art transformer-based FourCastNet
- Collaboration with NVIDIA, Caltech and others
- Forecasts global weather at 0.25° resolution
  - **Order of magnitude greater resolution** than state-of-the-art deep learning models
  - Forecasts wind speeds, precipitation and water vapor close to the skill of numerical weather prediction models up to 8 days
  - Produces a 24hr 100-member ensemble forecast in 7 seconds on a Perlmutter GPU node
  - Traditional NWP: 5 mins on *thousands of CPU nodes* for equivalent ensemble

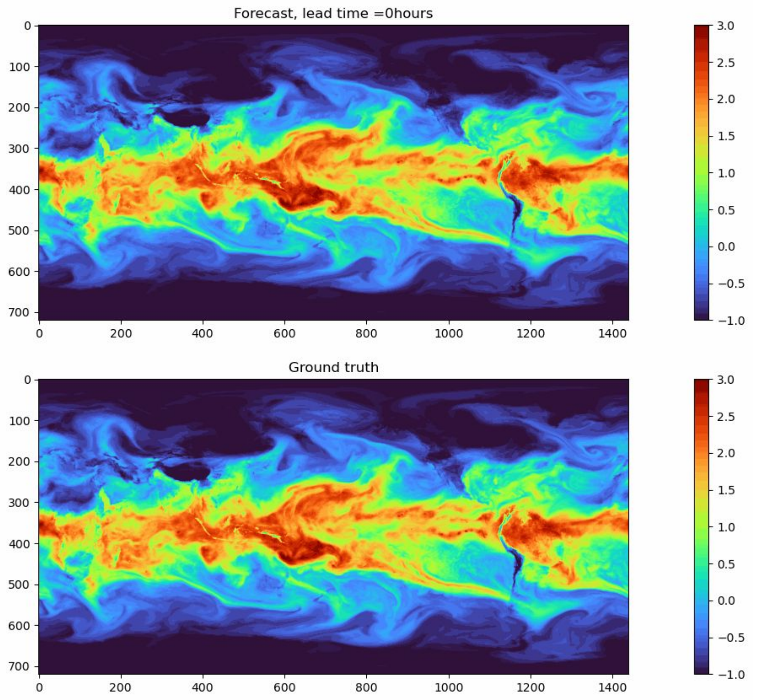

Data-driven forecast of an atmospheric river

Jaideep Pathak
former NERSC
Postdoc now NVIDIA

Shashank
Subramanian
NERSC Postdoc

Peter Harrington
NERSC ML
Engineer

BERKELEY LAB

Explore: Automated catalyst discovery

- GraphNNs to accelerate catalyst discovery for energy storage and climate change mitigation
- Collaboration with CMU and Facebook/Meta
- Created the largest catalysis dataset to date (OC20)
  - Challenge also in NeurIPS 2021 Competition
- Larger models achieve best performance on these large datasets - pushing to scale on Perlmutter

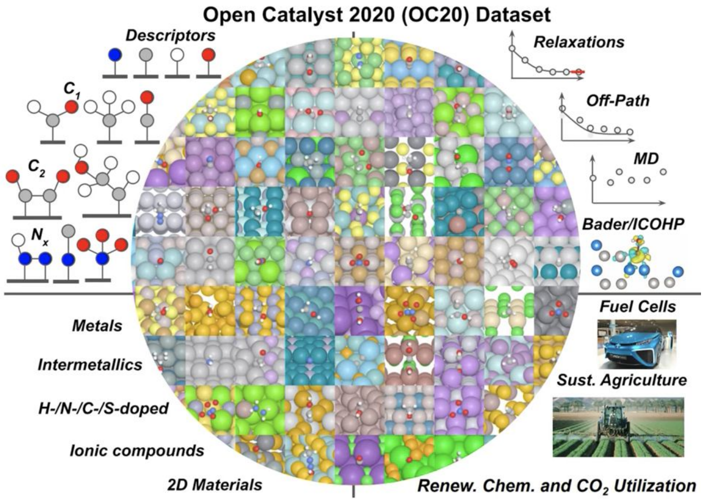

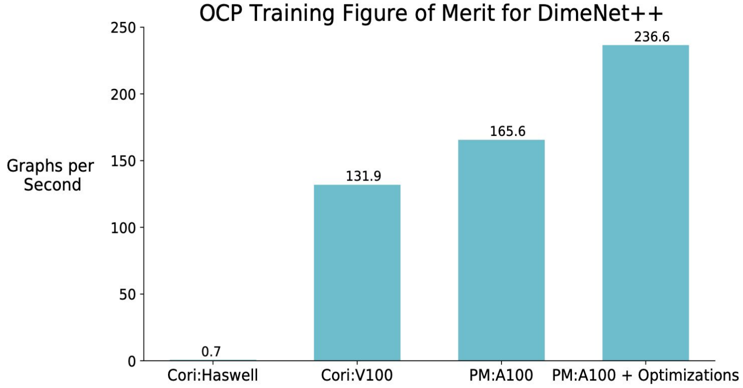

Performance comparison of Perlmutter (PM) with Cori CPU and GPU nodes. Optimizations carried out in collaboration with NVIDIA DevTechs

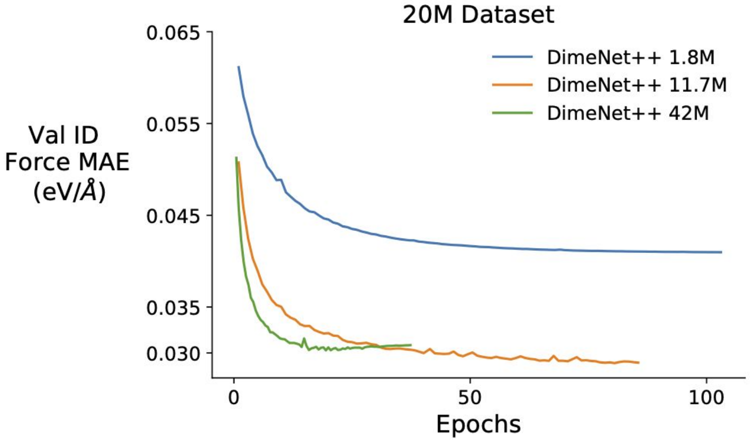

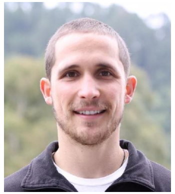

Brandon Wood
NERSC Postdoc now
Meta AI

Empowerment and training resources

The Deep Learning for Science School at Berkeley Lab https://dl4sci-school.lbl.gov/

- Lectures, demos, hands-on sessions, posters: 2019 in person (videos, slides, code)
- 2020 summer webinar series. Recorded talks: https://dl4sci-school.lbl.gov/agenda

The Deep Learning at Scale Tutorial

- Since 2018, and with NVIDIA in 2020/21
- 2021 was first training event to use Perlmutter Phase 1 with hands-on material for distributed training
- See the full SC21 material here and videos

NVIDIA AI for Science Bootcamp - Aug 25-26, 2022

Upcoming training for scientists interested in deploying DL
No experience required! Planning on using Perlmutter GPUs

Conclusions

- **AI for science requires supercomputing-scale capabilities:**
  - Hardware and software, application engagement and training
  - NERSC delivering this with Perlmutter
- **Usage of AI frameworks is growing. Need to:**
  - Provide optimized scalable software as well as flexibility for users
  - Allow for interactivity as well as automation
  - Utilize benchmarking for detailed performance tuning
- **Science AI projects reaching maturity and offer transformative potential**
  - Trend towards sophisticated science-specific architectures and scale
  - Examples running now on Perlmutter - much more to come
- **Future HPC systems: integrate AI into scientific workflows**
  - Optimizing for this will influence hardware design, system software, etc.

Thanks to all our staff and collaborators!

Steve Farrell, Wahid Bhimji, Mustafa Mustafa, Shashank Subramanian, Brandon Wood, …and many more from NERSC & Berkeley Lab.

Thorsten Kurth, Josh Romero, Jaideep Pathak (NVIDIA)

Questions?
Collaboration? Want to help?

Peter Harrington
pharrington@lbl.gov

Deep-learning@NERSC:
https://docs.nersc.gov/machinelearning/

BERKELEY LAB

NERSC

U.S. DEPARTMENT OF ENERGY | Office of Science