Scientific AI at Supercomputing Scale

Wahid Bhimji, Steve Farrell, Peter Harrington, Shashank Subramanian, Vinicius Mikuni
Data, AI and Analytics Services, NERSC

Featuring also
Mustafa Mustafa, Jaideep Pathak, Brandon Wood (former NERSC, Berkeley Lab) and others…

June 2023

NERSC: Mission HPC for the Dept. of Energy Office of Science

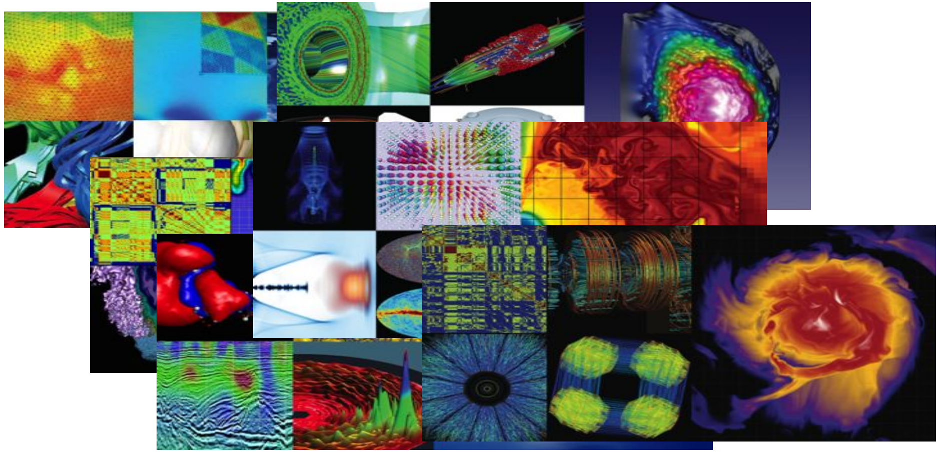

Large compute and data systems

- Perlmutter: ~7k A100 GPUs
- 128PB Community Filesystem ....

Broad science user base

- > 9,000 users,
- 1000 projects,

Outline

- AI for science is maturing and becoming transformative
- It benefits from supercomputing centres like NERSC
- Work still need across model development, approaches to scaling, compute systems and software
- Enabling scientific AI at scale requires developing cutting-edge applications and computing together

NERSC AI Strategy

Methods and Applications

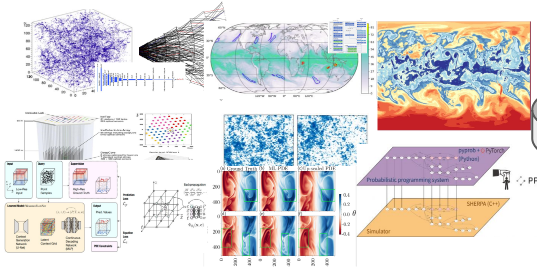

Deployment

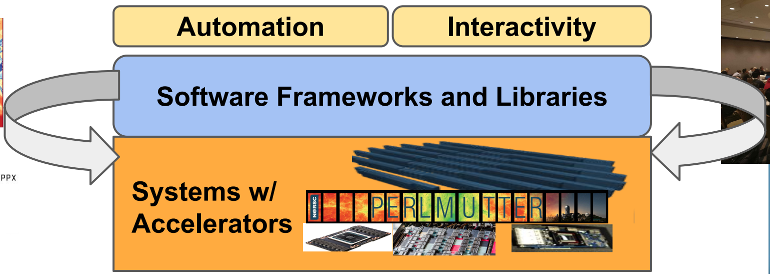

Empowerment

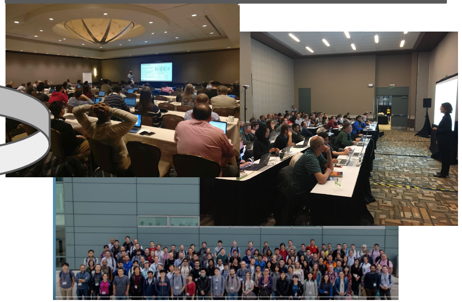

- ***Deploy*** optimized hardware and software systems
  - Currently Perlmutter >6000 A100 GPUs; Work with vendors for optimized AI software
  - >10x increase in number of users of DL frameworks from 2017 to 2021
  - Improve performance, e.g through benchmarking (e.g. MLPerf HPC))
- ***Apply*** *ML* for science using cutting-edge methods
  - “NESAP for Learning” application readiness program with postdocs, early access etc.
  - Other targeted engagements that push model development, scale and performance
  - Leverage lessons learned for all users
- ***Empower*** through seminars, training and schools
  - E.g. Deep Learning at Scale tutorial at Supercomputing (SC21 material here)

Transformative AI for new science - powered By Perlmutter

FourCastNet

Pathak et al. 2022 **arXiv:2202.11214**

Forecasts global weather at high-resolution. Hybrid data/ model parallel @ 4000 GPUs

First deep-learning with skill of numerical weather prediction

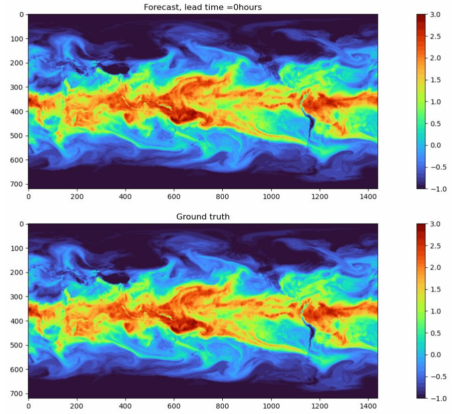

CatalysisDL

Chanussot et al. 2021 arXiv:2010.09990

Largest catalysis dataset (OC20 and OC22); Graph-parallel NN approaches and NeurIPS 2021 + 2022 Competitions

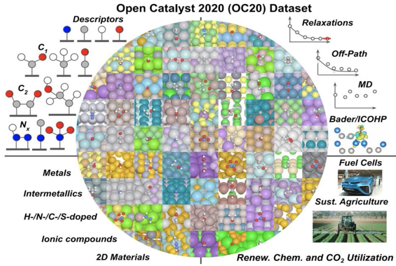

Pre-trained models now used with DFT - e.g. FineTuna; AdsorbML

Self-supervised sky surveys

Stein et. al. (2021) arXiv:2110.00023

Uncovered thousands of undiscovered strong-lenses

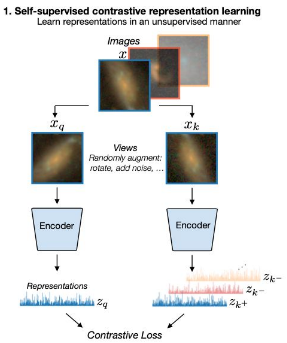

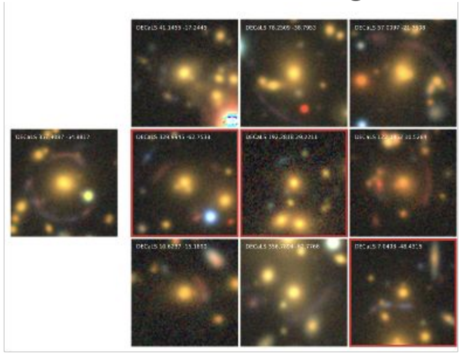

Unfolding for particle physics

H1 Collaboration ([...] Mikuni et. al.): recent press release

New ML approach extracts new physics insights.

Requires Perlmutter for 1000s of bootstrapping and UQ runs

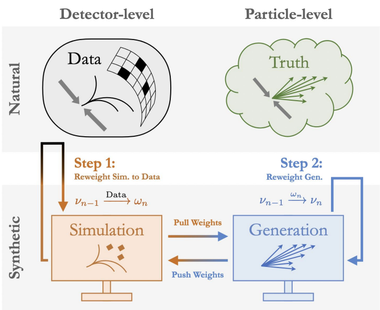

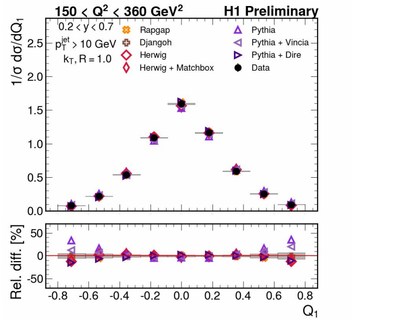

Deployment:
NERSC AI Systems, Workload and Software

Perlmutter: A Scientific AI Supercomputer

HPE/Cray Shasta system

Phase 1 (Dedicated May `21):

- 12 GPU cabinets with 4x NVIDIA A100 GPU nodes; Total >6000 GPUs
- 35 PB of All-Flash storage

Phase 2 (Integrated in 2022):

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

NVIDIA blog May 2021

See a growing scientific AI workload at NERSC

- Instrument user [python imports]
  - DL users >10x from 2017 to 2021
- Also track ML trends through 2-yearly survey
  - See similar trend in framework popularity

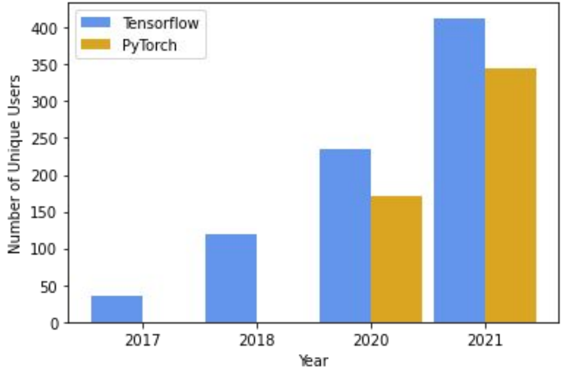

NERSC Survey: 2018 2020 2022

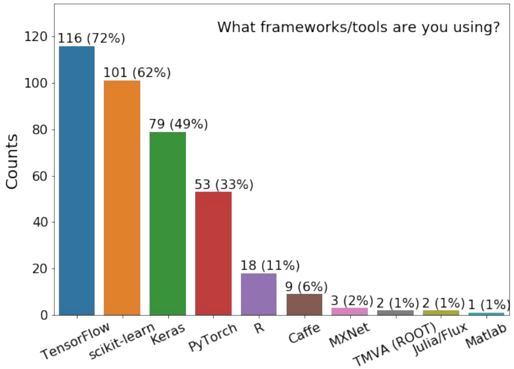

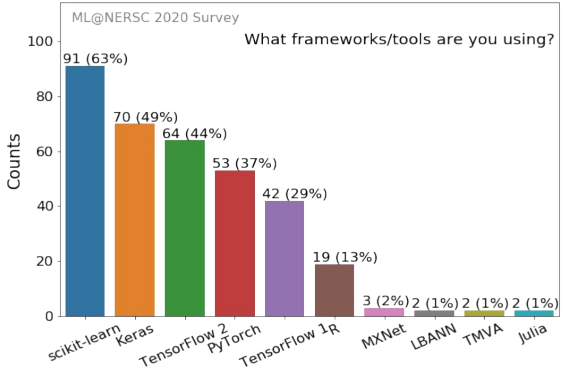

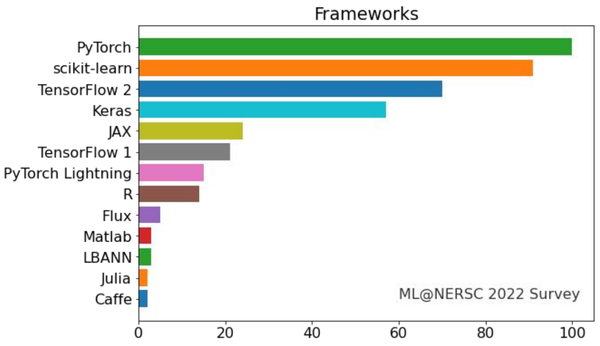

BERKELEY LAB

Current workload focuses on data analysis and image-like deep learning

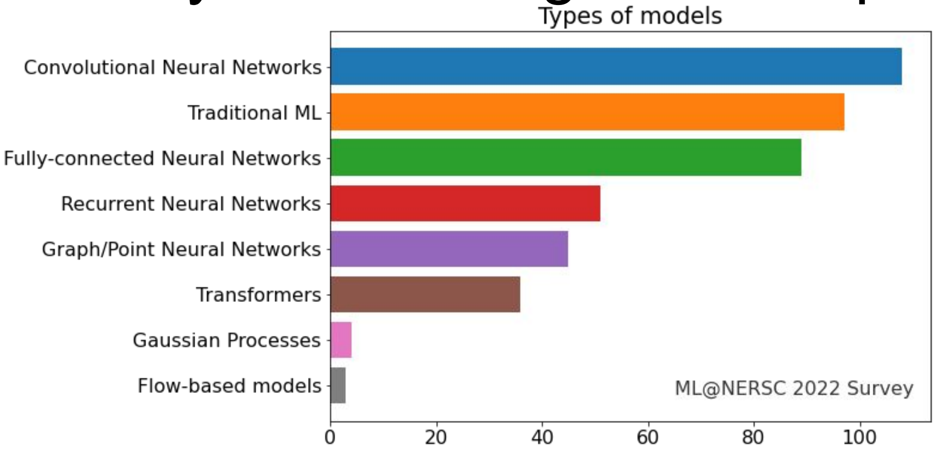

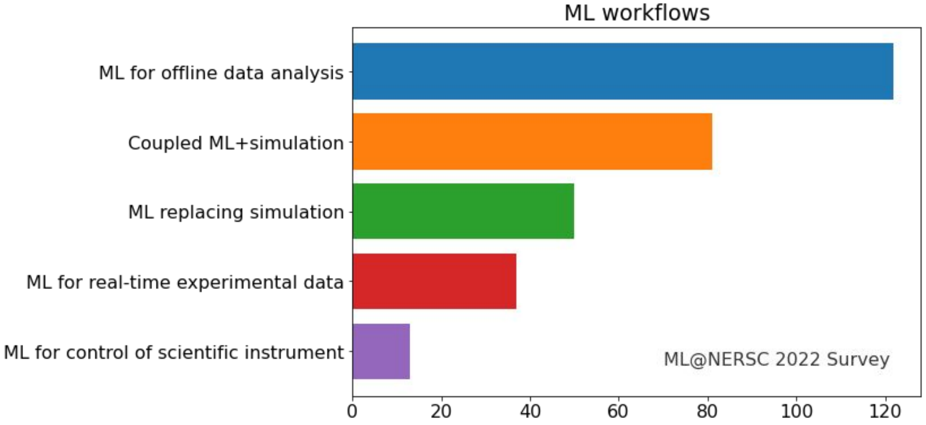

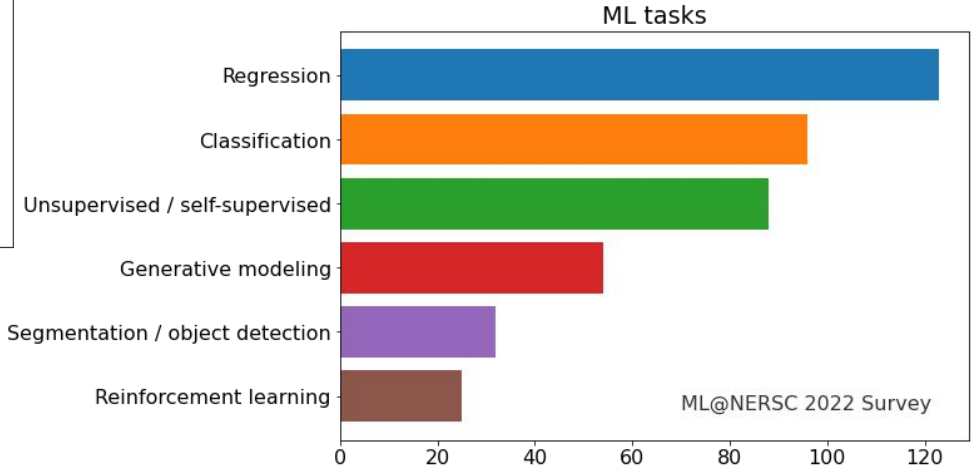

Need for AI at scale

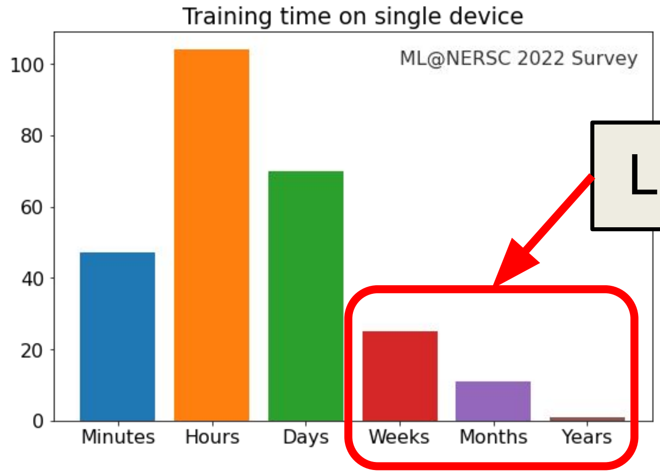

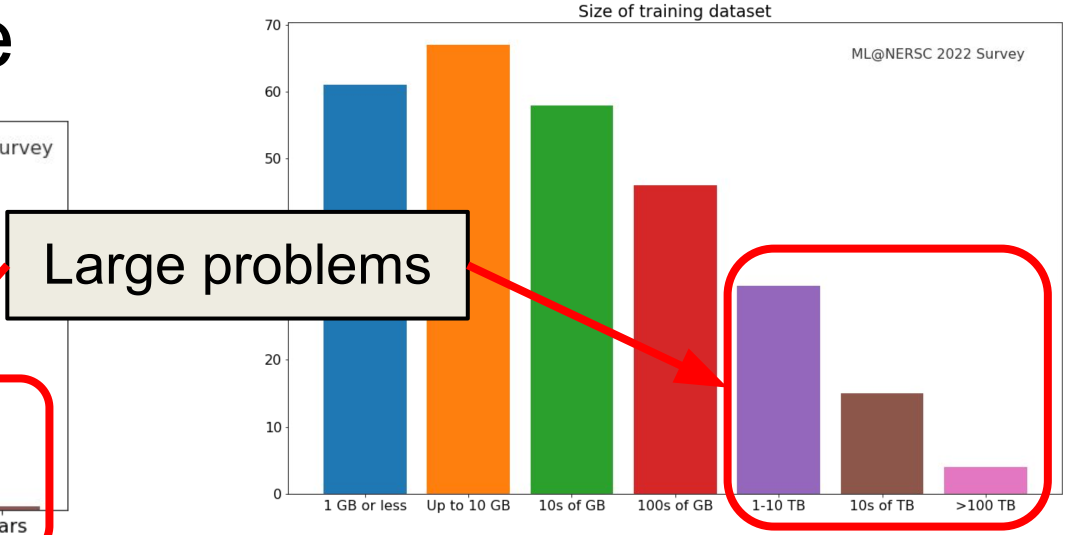

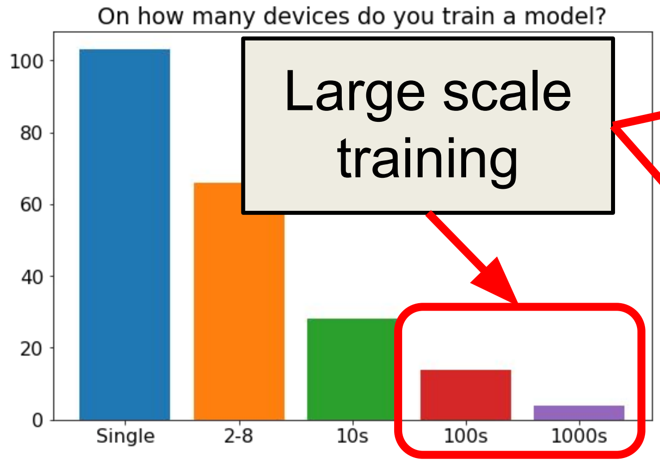

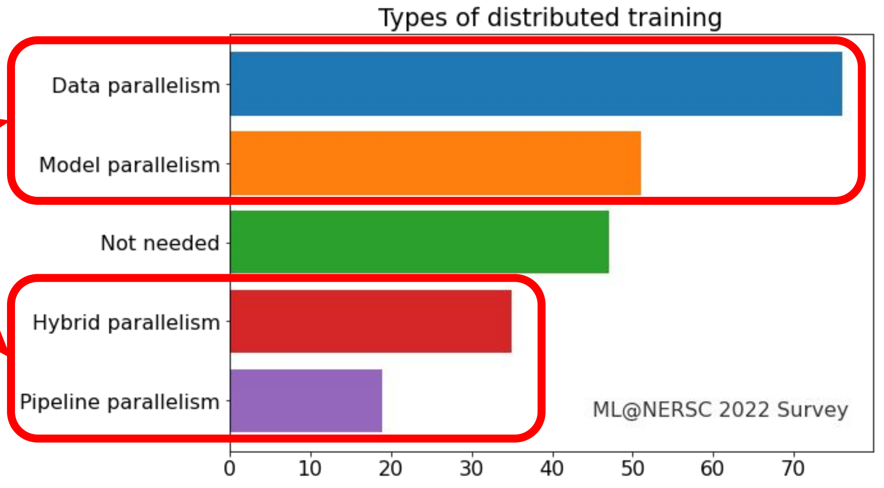

Scientists need performant and flexible software installations

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

- On Perlmutter we chose to deploy both compiled modules and NVIDIA's NGC containers
  - Container environment optimized for A100s and was crucial during deployment
  - Effectively debugged several deployment issues through close engagement with NVIDIA

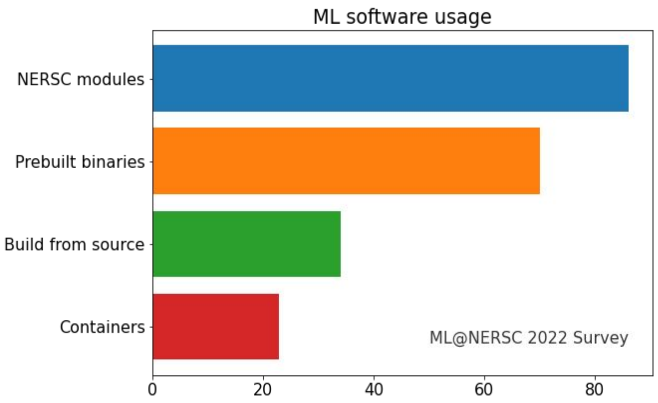

https://docs.nersc.gov/machinelearning/

Scientists need productive interfaces for experimentation

JupyterHub service provides a rich, interactive notebook ecosystem on Cori

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

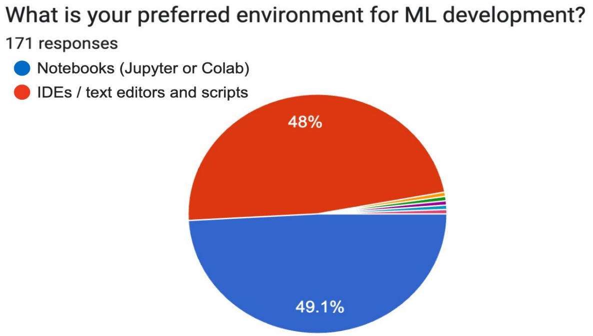

Users can run their deep learning workloads

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

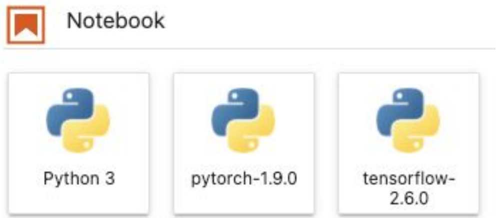

	Shared CPU Node	Shared GPU Node	Exclusive GPU Node	Exclusive Large Memory Node	Configurable GPU	Configurable DGX
Perlmutter	start		start		start	
Cori	start	start		start	start	start
Resources	Use a node shared with other users' notebooks but outside the batch queues.		Use your own node within a job allocation using defaults.		Use multiple compute nodes with specialized settings.	
Use Cases	Visualization and analytics that are not memory intensive and can run on just a few cores.		Visualization, analytics, machine learning that is compute or memory intensive but can be done on a single node.		Multi-node analytics jobs, jobs in reservations, custom project charging, and more.	

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

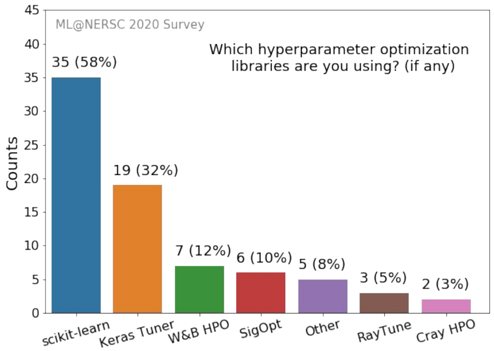

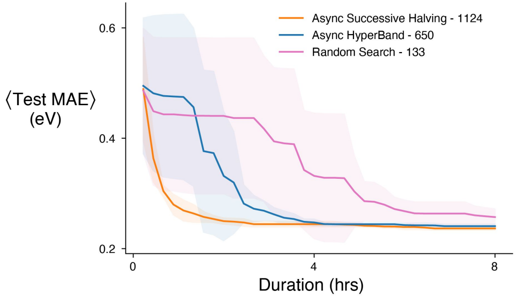

Multi-node RayTune HPO on Graph Neural Network models for catalysis applications (B. Wood et al.)

RAY tune Weights & Biases Determined AI SIGOPT OPTUNA

BERKELEY LAB

ML compute performance requires benchmarking and tuning

MLPerf™ is the industry standard benchmark for ML performance

For Science and Supercomputers: **MLPerf HPC** benchmark suite

ML Commons

- 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” (models/min) metrics
    - 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 at the SC21 MLHPC workshop
- **MLPerf HPC v2.0 presented at SC22**

https://mlcommons.org/en/get-involved/

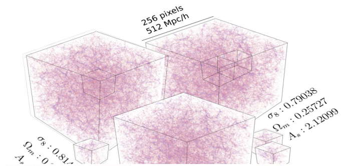

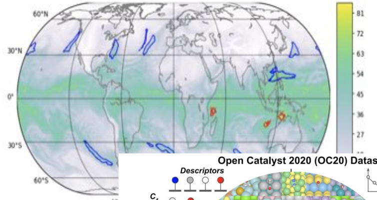

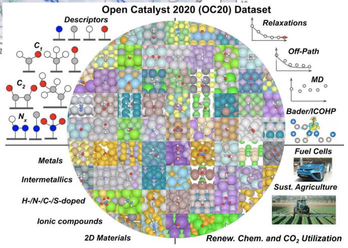

Applications and Empowerment:
Powered By Perlmutter

Transforming science with AI

Analyze

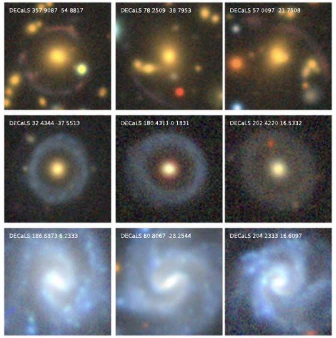

Hayat et al. 2021 [arXiv:2012.13083]

Accelerate

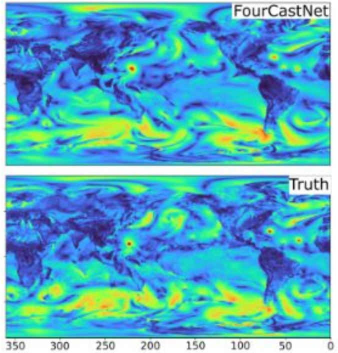

Pathak et al. 2022 [arXiv:2202.11214]

Automate

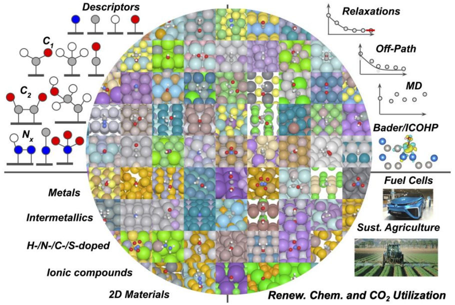

Chanussot et al. 2021
[arXiv:2010.09990]

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

Evolution of deep learning for science and *supercomputing*

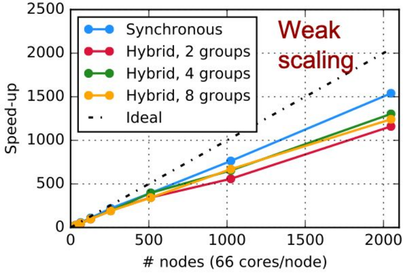

Some example projects:

- 2017 SC17 conference [Deep learning at 15PF]
- 2018 Gordon Bell Prize [Exascale DL for Climate Analytics]
- 2019 [Etalumis: bringing probabilistic programming to scientific simulators at scale]
- 2020 SC20 [MeshfreeFlowNet: a physics-constrained deep continuous space-time super-resolution framework]
- 2022 [FourCastNet: Accelerating Global High-Resolution Weather Forecasting using Adaptive Fourier Neural Operators]

This period showed a very rapid growth in

- Available Compute
  - 15 PetaFlops in SC17 -> 'Exascale' (half-precision) in SC18
- Sophistication of models and methods
- Availability of software
  - Custom hand-rolled Caffe/MPI SC17
  - Tensorflow/Horovod and Cray DL Plugin SC18
  - Pytorch DDP SC19

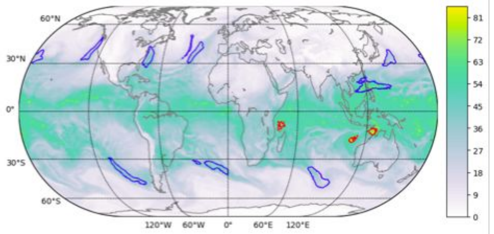

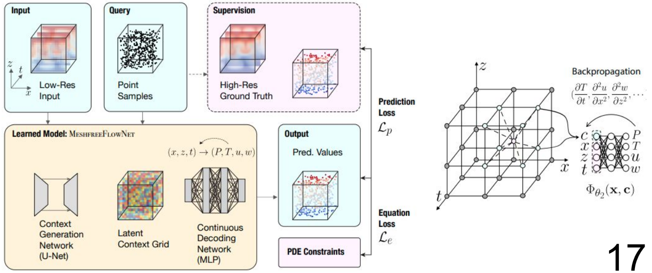

BERKELEY LAB

Analyze: Self-supervised sky surveys

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

- 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

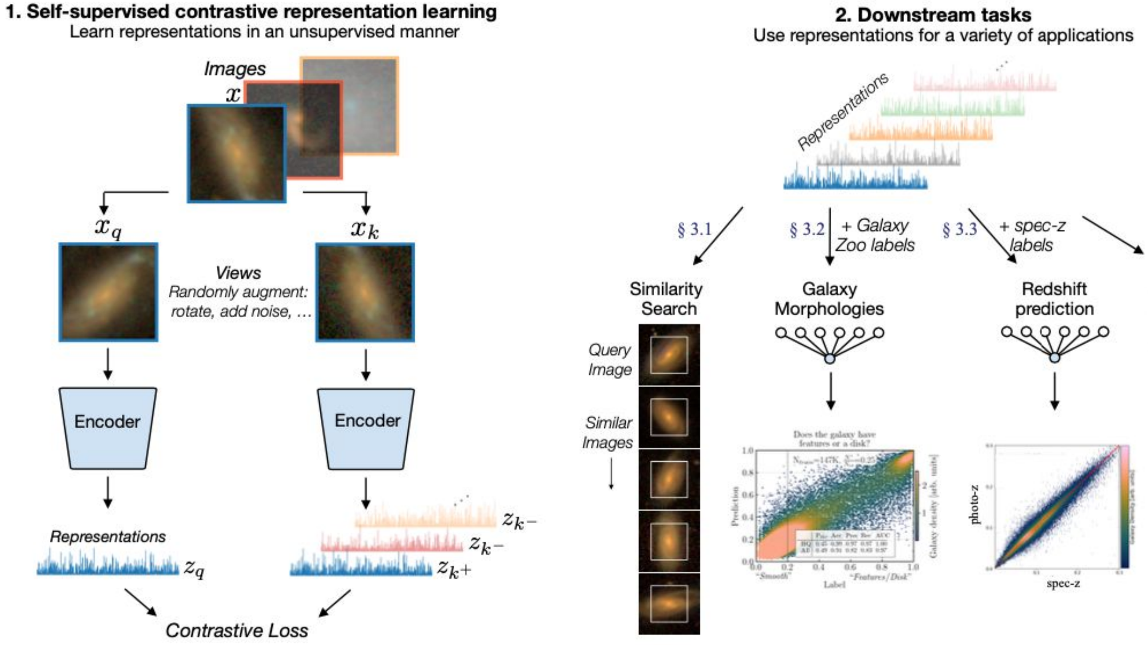

Detect & discover rare objects

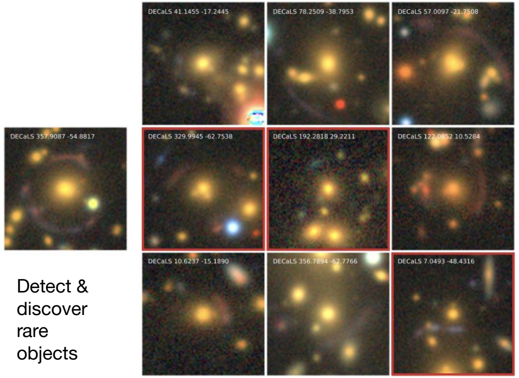

Peter Harrington
NERSC ML Engineer

Similarity search

- Given just a **single example**, instantly search for similar objects.

- Discover **new lenses or other phenomena** given just a few queries

Direction for future deep learning for science:

- **Community can benefit from multipurpose models trained on large-scale computing**

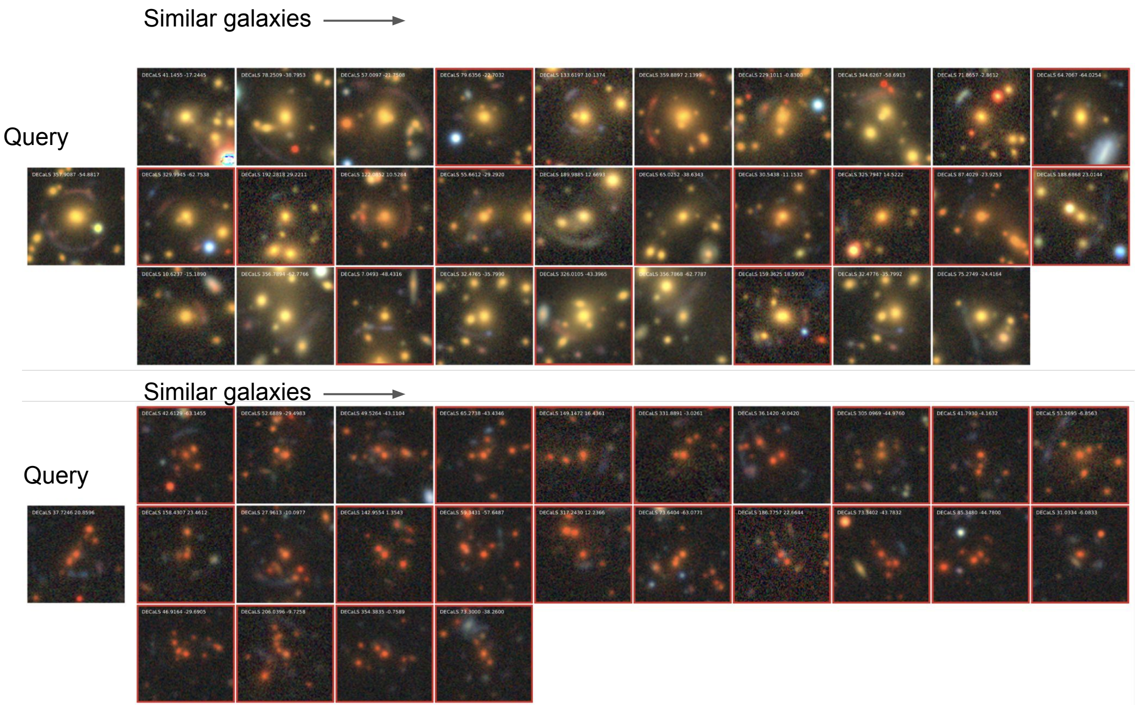

Try it out yourself:
share.streamlit.io/georgestein/galaxy_search

Accelerate: 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 “Fourier Neural Operator”
- 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

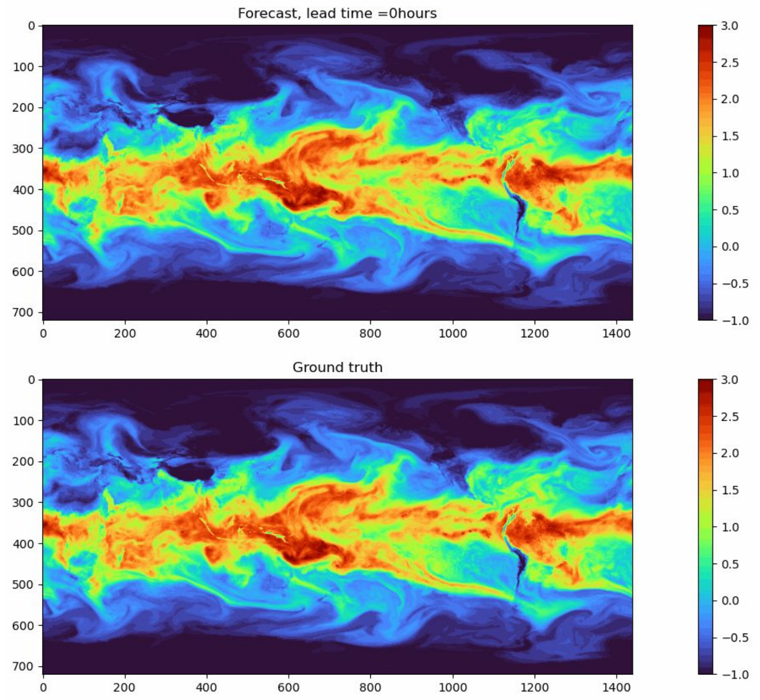

Data-driven forecast of an atmospheric river

Jaideep Pathak
former NERSC
Postdoc now NVIDIA

Peter Harrington
NERSC ML
Engineer

Shashank
Subramanian
NERSC Postdoc

BERKELEY LAB

FourCastNet: Large-compute scaling

Pathak et al. 2022
[arXiv:2202.11214](https://arxiv.org/abs/2202.11214)
Kurth et al. 2022
[arXiv:2208.05419](https://arxiv.org/abs/2208.05419)

Scales to e.g. 3808 GPUs on Perlmutter with model parallel on 4-gpus

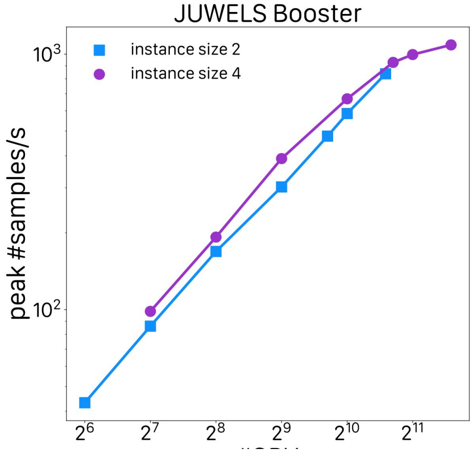

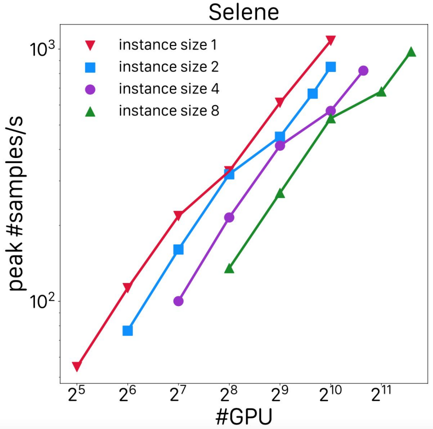

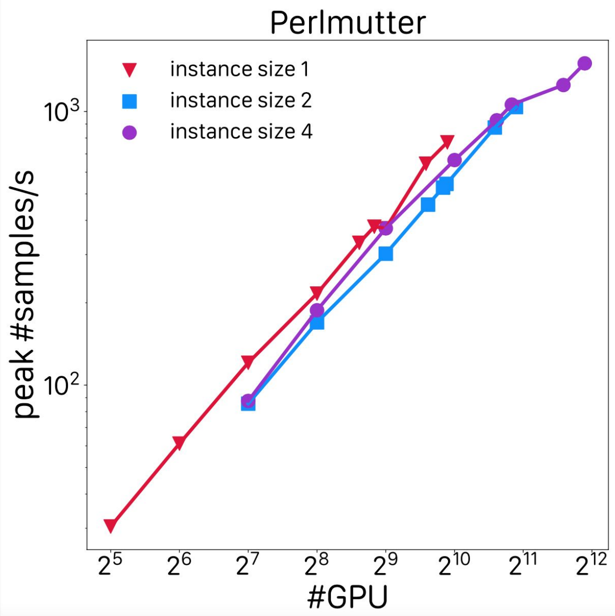

Train large models on ~1hr timescales compared to 40 hrs on 32 nodes or >~45days on a single GPU :

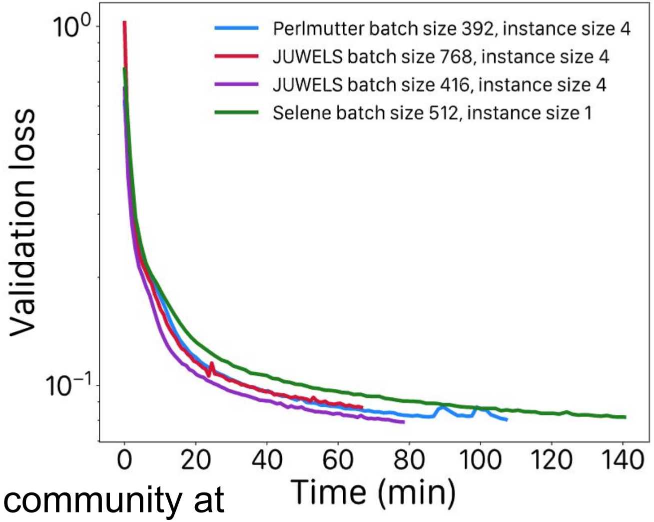

Model and weights made available to community at https://github.com/NVlabs/FourCastNet

Automate: discovering new catalysts

https://opencatalystproject.org/
Chanussot et al. 2021 arXiv:2010.09990

- GraphNNs to accelerate catalyst discovery for energy storage and climate change mitigation
- Collaboration with CMU and Facebook/Meta
- Largest catalysis datasets to date (OC20 and OC22)
  - Challenges in NeurIPS 2021 and 22
- Perlmutter helps push to larger better performing models
- Exploiting Graph-parallel NN approaches

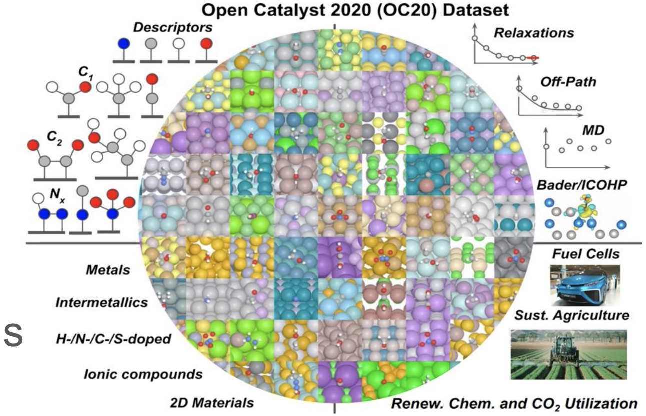

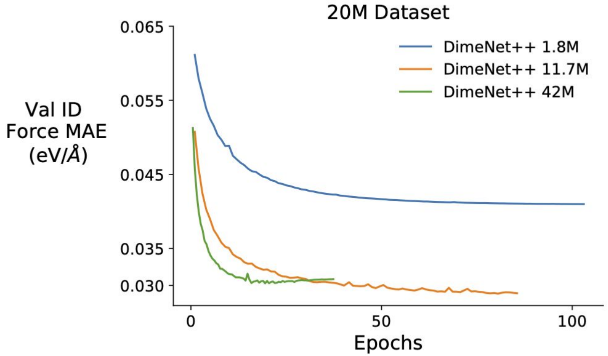

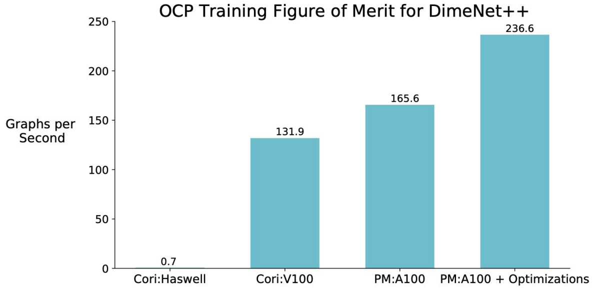

Performance comparison of Perlmutter (PM) with Cori CPU and GPU nodes.
Optimizations carried out in collaboration with NVIDIA DevTechs

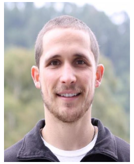

Brandon Wood
NERSC Postdoc now
Meta AI

- Public pre-trained models on OC20 now used by CMU group for 90% faster relaxation fine-tuned by DFT in active learning framework https://github.com/ulissigroup/finetuna

BERKELEY LAB

Unfolding for particle physics

H1 Collaboration ([...] Mikuni et. al. 2002 Phys. Rev. Lett. 128, 132002, 2022 Deep Inelastic Scattering (DIS) Conference. and recent press release

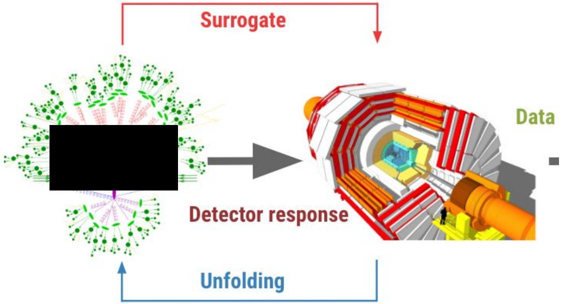

- “Unfolding” of fundamental particle interactions from observation in complex building-size experiments
- Collaboration with LBL Physics Division and H1 Collaboration
- Combines novel iterative ML approach OmniFold with GraphNN to extracts new physics insights
- Uses Perlmutter for 1000s of bootstrapping and UQ runs each using 128 GPUs for training

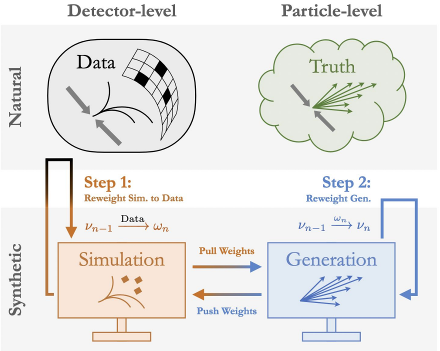

- **Other projects to replace full detector simulation** (expensive and not easily scalable)
  - **Using ML surrogate models** incorporating **diffusion generative models** for the first time in particle physics
  - **More info here:** arXiv:2206.11898

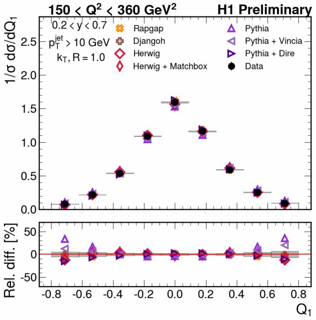

Vinicius Mikuni
NERSC Postdoc

BERKELEY LAB

Deep Learning on Supercomputers for Science 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 ***focussed on science and computing***.
  Recorded talks: https://dl4sci-school.lbl.gov/agenda

The *Deep Learning at Scale* Tutorial

- Run since 2018 with Cray, Intel and now OCLF and NVIDIA
- 2021 was powered by Perlmutter with **hands-on material for distributed training**
  - Full SC21 material here and videos
- Back in person for SC22 - material here

BERKELEY LAB

The future

NERSC-10

NERSC-10 will provide on-demand, dynamically composable, and resilient workflows across heterogeneous elements within NERSC and extending to the edge of experimental facilities and other user endpoints

Complexity and heterogeneity managed using complementary technologies

- **Programmable infrastructure**: avoid downfalls of one-size-fits-all, monolithic architecture
- **AI and automation**: sensible selection of default behaviours to reduce complexity for users

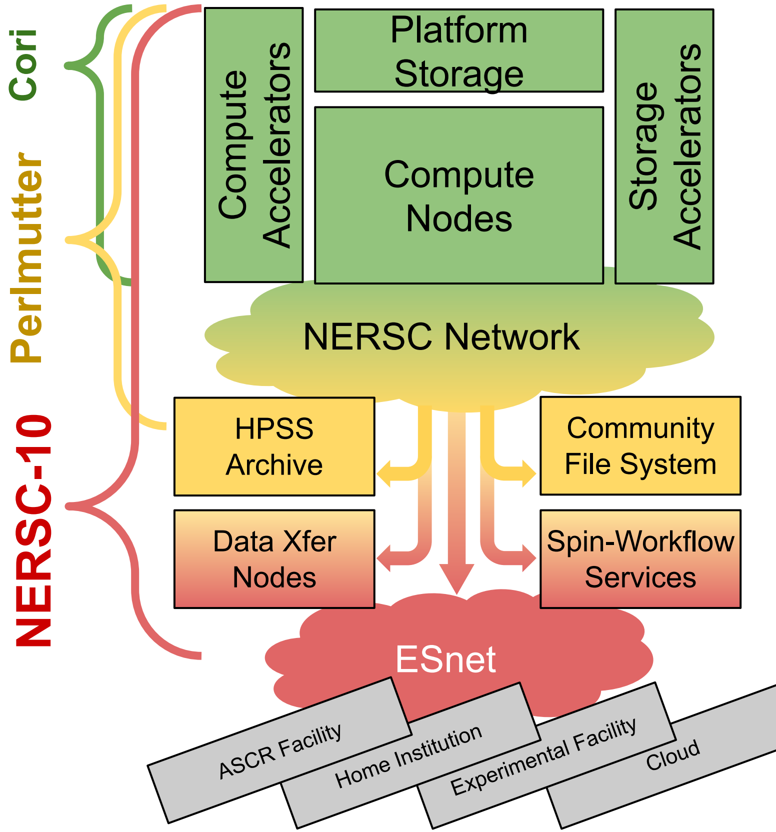

What we have

- Current best-in-class AI system - Perlmutter
- Flexible deep-learning frameworks: PyTorch; Tensorflow; and adding JAX
  - Optimised on Perlmutter through close collaboration with vendors
  - Detailed (and somewhat unique) tutorials with best practises on distributed scaling and performance profiling
- Productionized AI Models for Science in various domains
  - Trained at large-scale on Perlmutter
  - Workflows for uncertainty quantification and optimization using custom or open source tools

Because of this AI and Deep learning is being used *now* in production or near-production use cases for *multiple order of magnitude speedups* as well as *totally new science results*.

Numerous opportunities for R&D remain

- Hardware constraints limit size of AI models and/or push users to complex model and data-parallel strategies which require careful tuning
  - As well as compute, I/O can also be a bottleneck and data management tooling is limited
- Moving from single device to distributed AI resources requires user expertise and re-optimization
- AI models are expensive to train and specific to a domain and even to a sub-field/instrument and application
  - No *Foundation models* for science
- No standard frameworks or approaches for optimization, UQ and sharing
- Integration with HPC simulations or experimental data pipelines is limited, ad-hoc and very domain specific

With progress on these, AI for science can reach its *transformative potential*

5 year AI Directions

World-leading *open-science AI ecosystem* with

- **System** hardware and software that liberate scientists in application of AI models: including AI-acceleration, workflow and data management
- **Service platform** that offers:
  - Interactivity for large-scale model exploration
  - “Foundation” model hosting and retraining
    - Incorporating novel AI4Sci techniques
    - Accessible to AI novices and experts
  - Coupled AI, simulation and data pipelines
- **Science applications** with AI approaches that incorporate large-scale, science-informed, uncertainty-aware and transferable models
- ***Expertise, consulting and education***

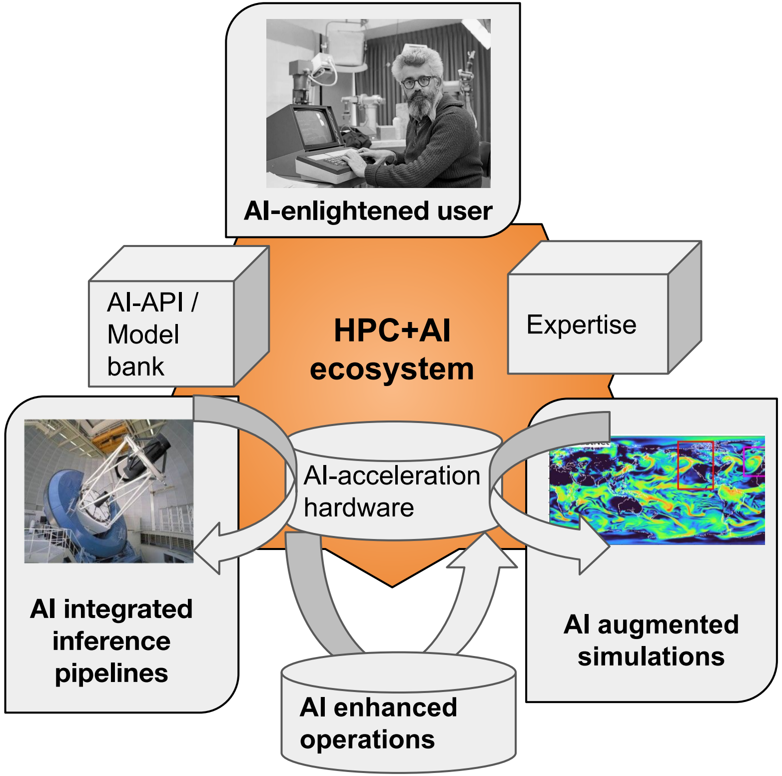

BERKELEY LAB

Ongoing AI Pathfinding at NERSC

Exploring *AI Acceleration and ecosystems*

- Evaluating the potential of **AI-focussed hardware** for science
  - Defining benchmarks and metrics for scientific ML - leading AI benchmarks on HPC systems: e.g. MLPerf HPC
  - Collating science experiences and deepening understanding of performance on AI-accelerated hardware - e.g. HPC DL Architectural requirements (PMBS 2021); MLPerf HPC analysis (MLHPC 21);
- Developing **novel applications** that fully exploit Perlmutter, current state-of-the-art AI system
  - Refresh NESAP applications - focus on production and scale
  - Joint projects with new LBL Machine Learning & Analytics CS-research Group, other LBL divisions and labs/universities
    - Scalable, transferable “Foundation” models for science
- Prototyping, evaluating and developing ***AI service platforms***
  - Vendor collaborations; market surveys; integrate open solutions

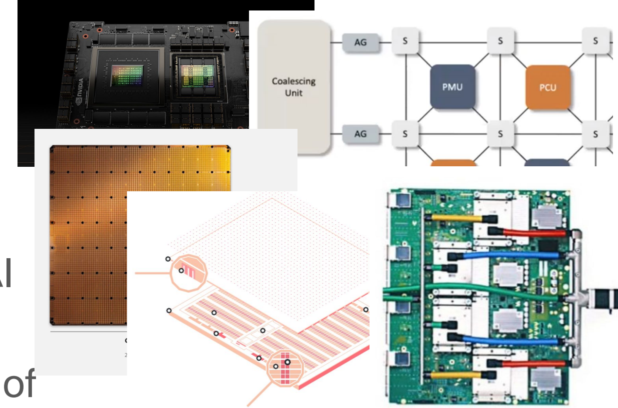

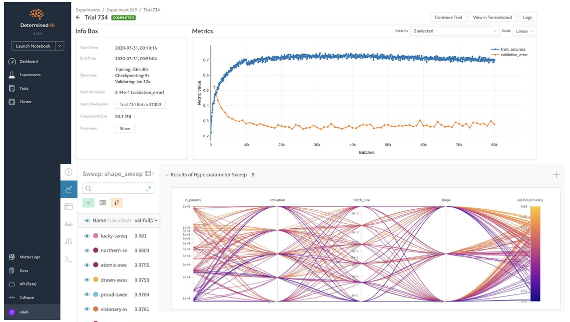

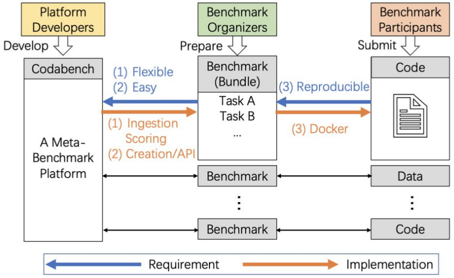

E.g New Fair Universe Project

Conclusions

- **Transformative AI for science should leverage supercomputing:**
  - Hardware, software, application engagement, and training
- **Usage of AI frameworks on HPC is growing. Need to:**
  - Provide optimized scalable software
  - Utilize benchmarking for detailed performance tuning
  - Allow for flexibility and interactivity as well as automation
- **Science AI projects reaching maturity with step-changes in performance**
  - Trends to training large models at scale that can be applied broadly
  - And utilizing sophisticated science-specific architectures
  - Recent results powered by Perlmutter - much more coming
- **Future HPC systems optimized to embed AI in workflows**
  - Hardware design, system software, features and tools

Questions?
Collaboration?

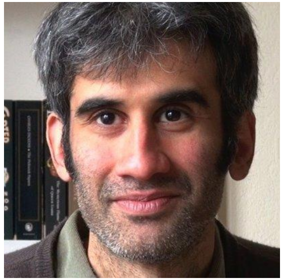

Wahid Bhimji
wbhimji@lbl.gov

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

BERKELEY LAB

NERSC

U.S. DEPARTMENT OF ENERGY | Office of Science