




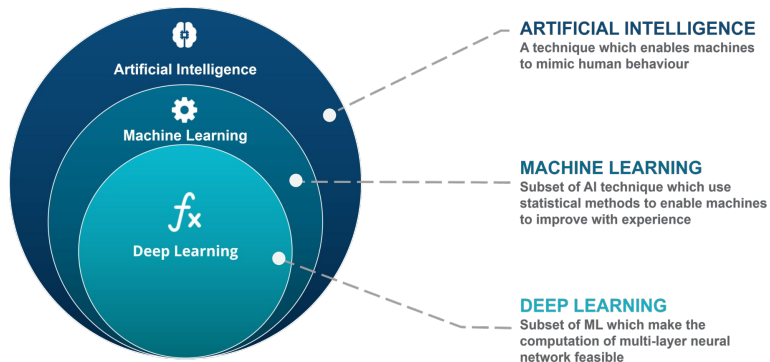
Deep Learning at NERSC



NERSC New User Training
September 8, 2023

Steven Farrell
Data, AI, and Analytics Services

The Deep Learning revolution



AI is transforming science




Across all domains

- Especially those with Big Data

Across *many* application areas

- Analyzing data better, faster
- Accelerating expensive simulations
- Control + design of complex systems

Embraced by the DOE and other funding agencies



**National Artificial Intelligence Research
Institutes**

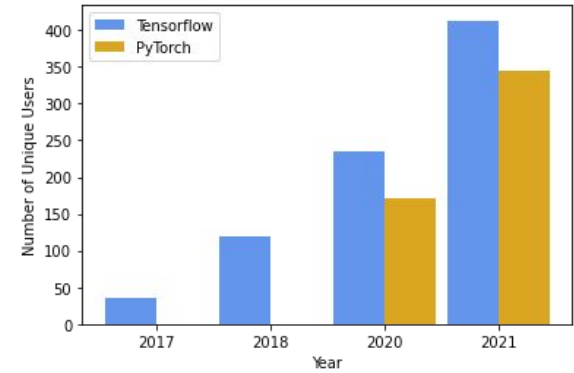


Growing scientific AI workload at NERSC

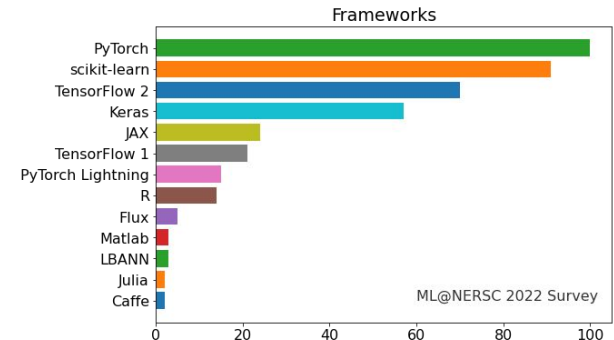

We track ML software usage

- Instrument user [python imports](#)
- DL users >10x from 2017 to 2021

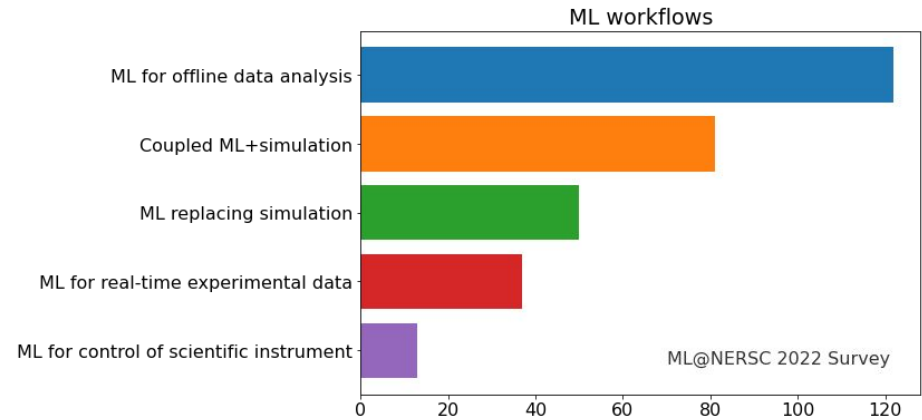

Also track ML trends through 2-yearly survey



Where do you run your ML models (include future plans)?





Scientific AI users



What is the level of maturity of ML in your research? (mark all that apply to your projects)

174 responses



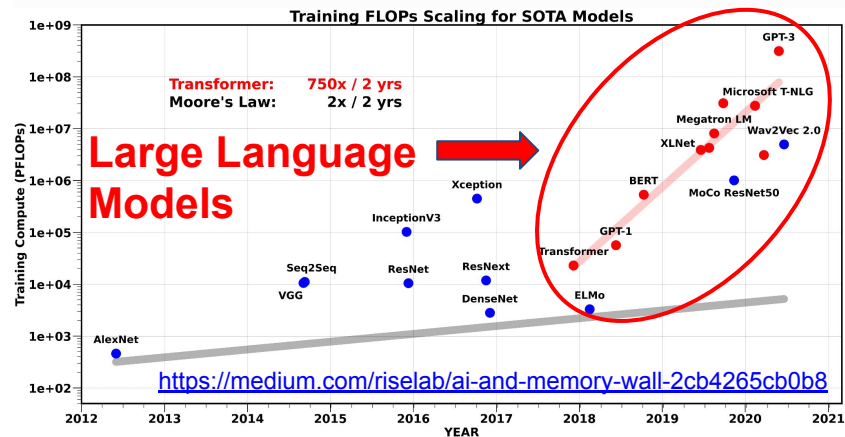
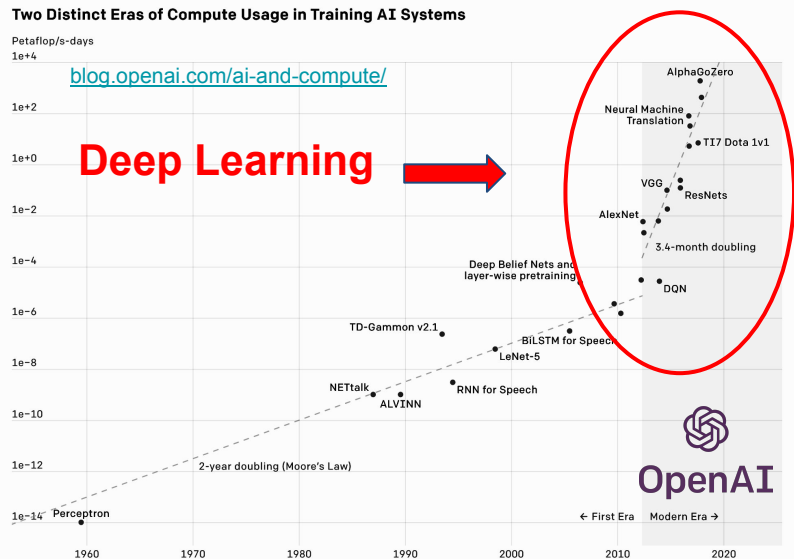
The need for HPC

Growing computational cost of training AI models



- bigger datasets + models, more complexity

Researchers need large scale resources


- Rapid iteration, reduce time to discovery




Need for AI at scale



Large problems




Large scale training




ML@NERSC 2022 Survey

Deep Learning parallelization strategies




Data Parallelism

Distribute input samples. Distribute network structure (layers).



Model (tensor) Parallelism

Distribute network structure (layers).



Layer Pipelining

Partition by layer.

Fig. credit: [arXiv:1802.09941](https://arxiv.org/abs/1802.09941)

Hybrid parallelism example: [Megatron-Turing NLG 530B](#)

Current NERSC AI Strategy

Deployment

Automation

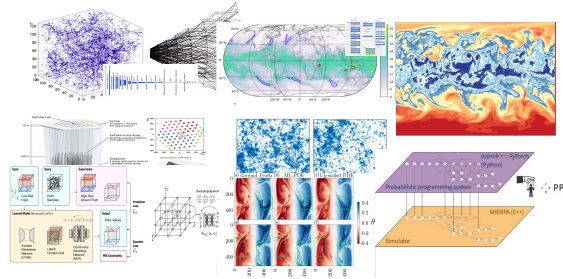
Interactivity

Software Frameworks and Libraries

Systems w/
Accelerators



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

Deep Learning on Perlmutter: Software stack and best practices



Perlmutter deep learning software stack overview

General strategy:

- Provide functional, performant installations of the most popular frameworks and libraries
- Enable flexibility for users to customize and deploy their own solutions

Frameworks:




Distributed training libraries:

- Horovod
- PyTorch distributed
- NCCL, MPI

Productive tools and services:

- Jupyter, Shifter



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

How to use the Perlmutter DL software stack

We have modules you can load which contain python and DL libraries:

```
module load tensorflow/2.9.0
```

```
module load pytorch/2.0.1
```

Check which software versions are available with:

```
module spider pytorch
```

You can install your own packages on top to customize:

```
pip install --user MY-PACKAGE
```

Or, clone a conda environment from our modules:

```
conda create -n my-env --clone /path/to/module/installation
```

Or, create custom conda environments from scratch:

```
conda create -n my-env MY-PACKAGES
```

More on how to customize your setup can be found in the docs ([TensorFlow](#), [PyTorch](#)).

Containerized DL: using Shifter on Perlmutter

NERSC currently supports [containers with Perlmutter via Shifter](#)

- Easy, performant: our top500 entry used a container!

To see images currently available:

```
shifterimg images | grep pytorch
```

To pull desired docker images onto Perlmutter:

```
shifterimg pull <dockerhub_image_tag>
```

To use interactively:

```
shifter --module gpu --image=nvcr.io/nvidia/pytorch:22.05-py3
```

Use Slurm image shifter options for best performance in batch jobs:

```
#SBATCH --image=nersc/pytorch:ngc-22.05_v1
#SBATCH --module=gpu,nccl-2.15
srun shifter python my_python_script.py
```

Coming soon: Podman!



Best Practices for DL + Shifter on Perlmutter

NVIDIA provides [containers optimized for deep learning on GPUs](#) with

- Pytorch or TensorFlow+Horovod
- Optimized drivers, CUDA, NCCL, cuDNN, etc
- Many different versions available



We also provide [images](#) based on NVIDIA's, which have a few useful extras

You can also build your own custom containers (easy to build on top of NVIDIA's)

Notes

- [Customization](#): from inside the container, do `pip install --user MY-PACKAGE` (make sure to set `$PYTHONUSERBASE` to a custom path for the desired container)
- NVIDIA NGC containers use OpenMPI, which requires specific options if you require MPI. Instructions: <https://docs.nersc.gov/development/shifter/how-to-use/#shifter-mpich-module>

Distributed Training Tools

Framework built-in


- PyTorch DistributedDataParallel (DDP)
- TensorFlow Distribution Strategies

Other popular libraries

- **Horovod**: MPI+NCCL, easy to use, [examples](#)
- **Lightning**: DDP + convenient features
- **DeepSpeed**: ZeRO optimizations, 3D parallelism
- **Ray**: DDP + HPO
- **LBANN**: multi-level parallelism, ensemble learning, etc., [docs](#)

Communication backends

- NCCL is the backend of choice for GPU nodes on Perlmutter
- The NCCL OFI plugin (from AWS) enables RDMA performance on the libfabric-based Perlmutter Slingshot network (see our docs)



General guidelines for deep learning at NERSC

NERSC documentation: <https://docs.nersc.gov/analytics/machinelearning/overview/>

Use our provided modules/containers if appropriate

- They have the recommended builds and libraries tested for functionality and performance
- We can track usage which informs our software support strategy


For developing and testing your ML workflows

- Use interactive QOS or Jupyter for on-demand compute resources
- Visualize your models and results with TensorBoard or Weights & Biases

For performance tuning

- Check cpu/gpu utilization to indicate bottlenecks (e.g. with top, nvidia-smi)
- Data pipeline is the most common source of bottlenecks
 - Use framework-recommended APIs/formats for data loading
 - Use multi-threaded data loaders and stage data if possible
- Profile your code, e.g. with Nvidia Nsight Systems or TensorBoard Profiler

Deep Learning on Perlmutter: Workflow tools



Jupyter for deep learning

JupyterHub service provides a rich, interactive notebook ecosystem on Cori

- Very popular service with hundreds of users
- 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
- [using their own custom kernels](#)



Notebook



	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	
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 s	
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.	

TensorBoard at NERSC


TensorBoard is the most popular tool for visualizing and monitoring DL experiments, widely adopted by TensorFlow and PyTorch communities.

We recommend running TensorBoard in Jupyter using [nersc-tensorboard helper module](#).

```
import nersc_tensorboard_helper
%load_ext tensorboard
%tensorboard --logdir YOURLOGDIR --port 0
```

then get an address to your TensorBoard GUI:

```
nersc_tensorboard_helper.tb_address()
```



Model selection/tuning are critical for getting the most out of deep learning

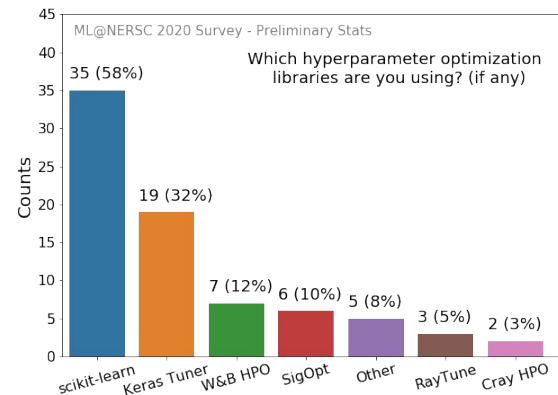
- Many methods and libraries exist for tuning your model hyper-parameters
- Usually very computationally expensive because you need to train many models
=> Good for large HPC resources

Helpers / examples


- [W&B template \(new\)](#)
- [Ray cluster helper \(new\)](#)

Users can use whatever tools work best for them

- Ask us for help if needed!



Outreach & additional resources



Training events

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

- Comprehensive program with lectures, demos, hands-on sessions, posters
- 2019 material (videos, slides, code) online: <https://sites.google.com/lbl.gov/dl4sci2019>
- 2020 webinar series material: <https://dl4sci-school.lbl.gov/agenda>

The Deep Learning at Scale Tutorial

- Jointly organized with NVIDIA (+ previously Cray, ORNL)
- Presented at SC18-22, ECP Annual 2019, ISC19
- Detailed lectures + hands-on material covering distributed training, scaling, profiling, and optimization on Perlmutter
- [See the full SC22 material here](#)
- *Accepted for SC23 (w/ more model parallelism)*

NVIDIA AI for Science Bootcamp

- [2022 event](#)
- [2023 event](#) (Oct 18, *apply now!!*)

NERSC Data Seminar Series:

- <https://github.com/NERSC/data-seminars>



Conclusions

Deep learning for science is here and growing

- Powerful capabilities
- Enthusiastic community
- Increasing HPC workloads

Perlmutter has a productive, performant software stack for deep learning

- Optimized frameworks and solutions for small to large scale DL workloads
- Support for productive workflows (Jupyter, HPO)

Join the [NERSC Users Slack](#)

Take the [ML@NERSC Survey](#)





Thank You and
Welcome to
NERSC!



Data-parallel training considerations

Weak scaling: converge faster by taking fewer, bigger, faster steps

- i.e., more GPUs, larger batch sizes, larger learning rates



Upper: 3 SGD steps w. learning-rate = η

Lower: 1 SGD step w. learning-rate = $3 * \eta$

Caveat: for stability & convergence, requires tuning

- Warm-up+scale learning rate, adaptive optimizers, etc
- See our [SC21 “Deep Learning at Scale” tutorial](#) for more tips

Model parallelism

Why you might try model parallelism

- to fit larger models
- for speedup (results may vary)

Generally, you can combine multiple types of parallelism

- mention some example like nvidia megatron
- <https://www.microsoft.com/en-us/research/blog/using-deepspeed-and-megatron-to-train-megatron-turing-nlg-530b-the-worlds-largest-and-most-powerful-generative-language-model/>

Disclaimer: not much content today

- mention some tools
- mention plans for SC23 tutorial

Guidelines - TensorFlow distributed training

TensorFlow at NERSC docs:

<https://docs.nersc.gov/analytics/machinelearning/tensorflow/>



TensorFlow

For distributed training, we recommend using Horovod

- Easy to use and launch with SLURM
- Can use MPI and NCCL as appropriate
- Horovod examples:

<https://github.com/horovod/horovod/tree/master/examples>



TensorFlow has some nice built-in profiling capabilities

- TF profiler in TF 2: <https://www.tensorflow.org/guide/profiler>

Guidelines - PyTorch distributed training

PyTorch at NERSC docs:

<https://docs.nersc.gov/analytics/machinelearning/pytorch/>



For distributed training, use PyTorch's DistributedDataParallel

- Simple model wrapper, native to Pytorch
- Works on CPU and GPU
- Highly optimized for distributed GPU training
- Docs:

https://pytorch.org/tutorials/intermediate/ddp_tutorial.html

Distributed backends

- On Perlmutter, use the NCCL backend for optimized GPU communication



What are we working on now?

Inference serving

Platforms / ecosystems for AI workflows and MLOps

Podman-hpc for AI

Large scale distributed training, HPO, inference from Jupyter notebooks

Deep Learning is transforming science

It can enhance various scientific workflows




- Analysis of large, complex datasets
- Accelerating expensive simulations

Adoption is on the rise in the science communities

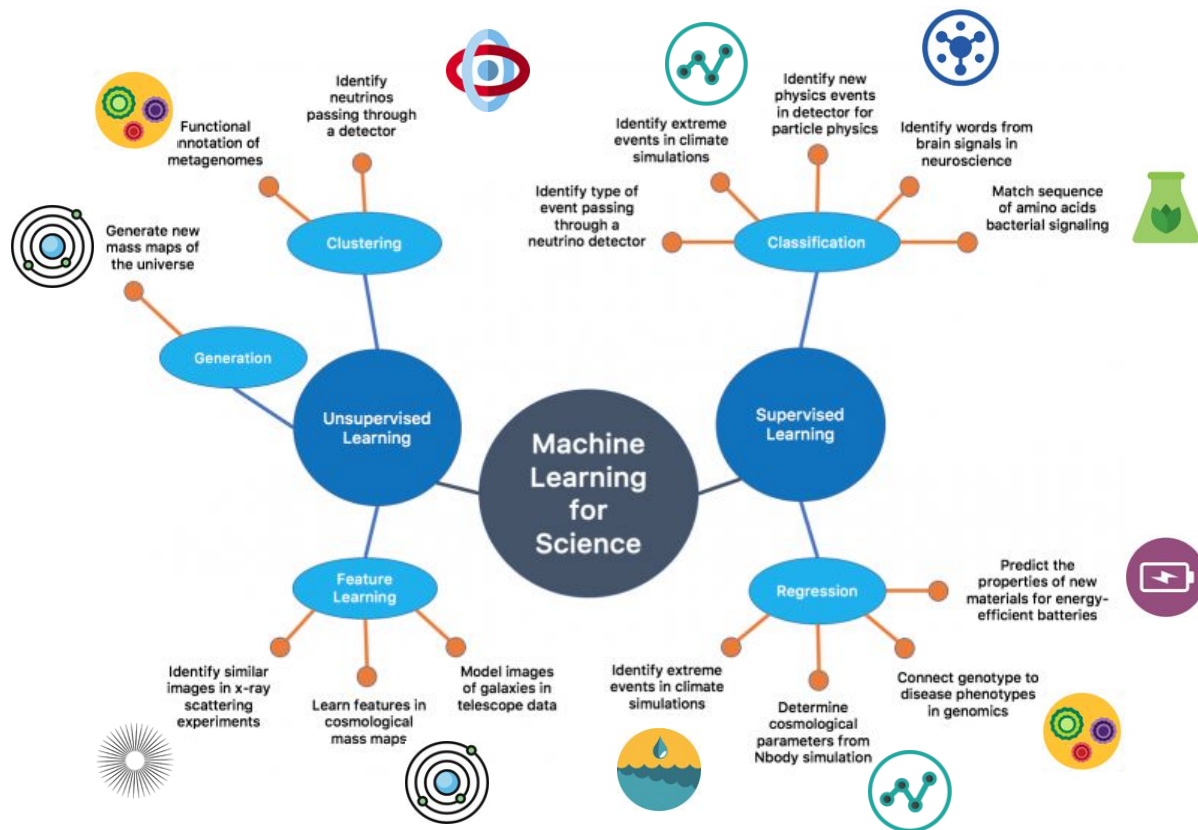
- Rapid growth in ML+science conferences
- Recognition of AI achievements:
2018 Turing Award; 2018, 2020 Gordon Bell prizes
- HPC centers awarding allocations for AI,
optimizing next-gen systems for AI

The DOE is investing heavily in AI for science


- Funding calls from ASCR (and other funding agencies)
- Popular, enthusiastic AI4Science town hall series, [300 page report](#)



Scientific ML: endless possibilities!




More complex tasks, bigger models, more compute




Credit: NVIDIA

At what scale do you train your models? (include current and future plans).

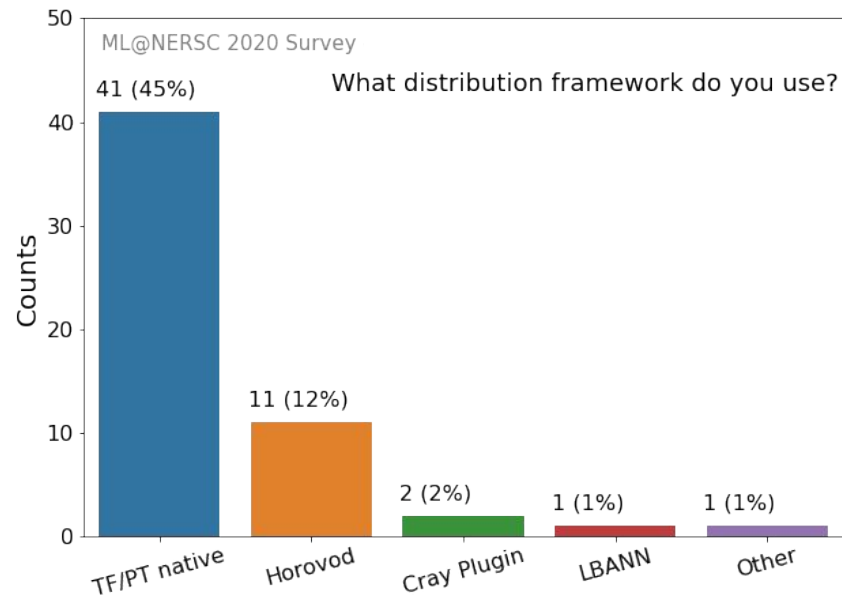


Models get bigger and more compute intensive as they tackle more complex tasks

Deep Learning parallelization strategies



Data parallelism is the most common strategy in practice, especially for inter-node scaling.



TensorFlow and PyTorch support data and intra-node pipeline parallelism natively. Horovod is the leading non-native distribution framework. All support MPI and/or NCCL backends.