

# The PATH to NAIRR: Heterogeneous Training of Machine Learning Models in OSPool and NAIRR

Ian Ross  
HTC26 - June 10, 2026

This project is supported by the National Science Foundation under Cooperative Agreements OAC-2030508. Any opinions, findings, conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

# Observation - AI is here and has altered scientific processes.


AI is **changing the way that science is being done.**

I won't talk (here...) about

- Modalities of “AI-driven science”
- Hype vs reality
- Safety
- Best practices
- Agent-driven software development
  - And what it means to education and facilitation

Instead, this talk focuses on a experiment showing that **distributed high throughput computing is applicable and powerful in AI-driven science, training 30+ complete models, with >1,200 epochs of training and almost 20,000 GPU hours**

THE AI REVOLUTION: PANELS OF UNFULFILLED POTENTIAL



Nano Banana 2 prompt: "Create a funny comic of an AI (robot) who isn't \_quite\_ living up to all the hype just yet"

# Hypothesis

There is **value** in applying the principles of distributed high throughput computing to certain **machine learning tasks**:

- **Ensembles** of models (mixture-of-experts, targeted protein/physics/simulations)
- **Ablation studies** of small-to-medium sized models
- **Opportunistically shaped** work (“give me any GPU for any time”)
- **Inference**
- **Automated workflows**
- **Teaching**

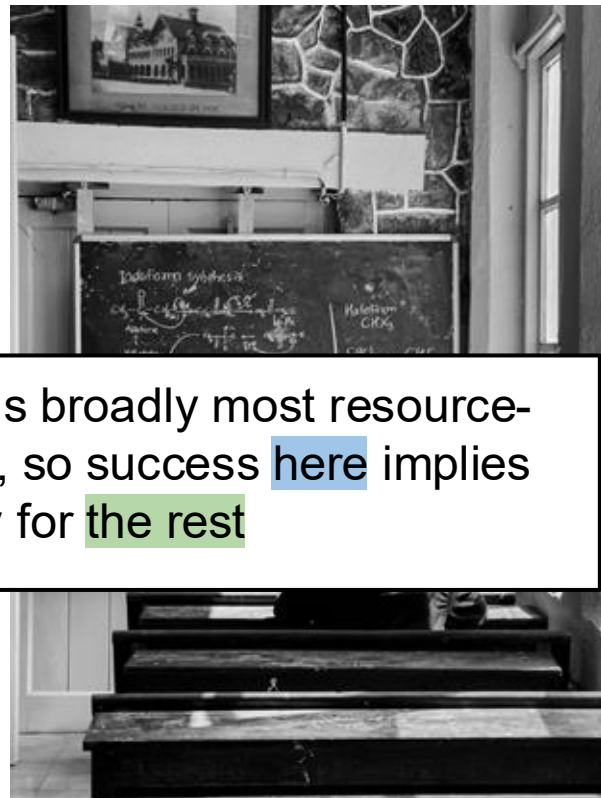


# Hypothesis

There is **value** in applying the principles of distributed high throughput computing to certain **machine learning tasks**:

- **Ensembles** of models (mixture-of-experts, targeted protein/physics/simulations)
- **Ablation studies** of small-to-medium sized models
- **Opportunistically shaped** work (“give me any GPU for any time”)
- **Inference**
- **Automated workflows**
- **Teaching**

Training is broadly most resource-intensive, so success **here** implies suitability for **the rest**



# Experiment objectives



Photo by [Drew](#) on [Unsplash](#)

Run a suite of training tasks to:


🔬 Characterize the impact of training models across heterogeneous resources

♻️ (Translational Computer Science) Test and advance PATH Access Point capabilities to reduce the barrier of entry and support ML workloads in distributed and heterogeneous environments

# Methodology

Covered in [more detail last year](#) but briefly:

- Train variants of a “known” scientific model (Gitter Lab’s METL protein language model), intentionally moving to new sites at epoch boundaries
  - Leverage NAIRR resources (via HTCondor Annex) in addition to OSPool for “heterogeneously trained” approach
  - Use CHTC, sticking to single host or single GPU variety for baseline
- Using validation dataset, compare:
  - Training behavior (test loss, convergence behavior)
  - Science power (Pearson correlations between various energy/shape predictions, compared to Rosetta simulations)



# The impact of NAIRR...

The targeted language model requires ~60GB of VRAM. OSPool contributions exist at this size, but researcher needs continue to grow. We wanted to make sure a path exists if (when?) some users outgrow OSPool capabilities.

Enter the National AI Research Resource (NAIRR):

“The National Artificial Intelligence Research Resource (NAIRR) is a concept for a shared national research infrastructure to bridge this gap by connecting U.S. researchers to responsible and trustworthy Artificial Intelligence (AI) resources, as well as the needed computational, data, software, training, and educational resources to advance research, discovery, and innovation.” [From nairrpilot.org](https://nairrpilot.org)

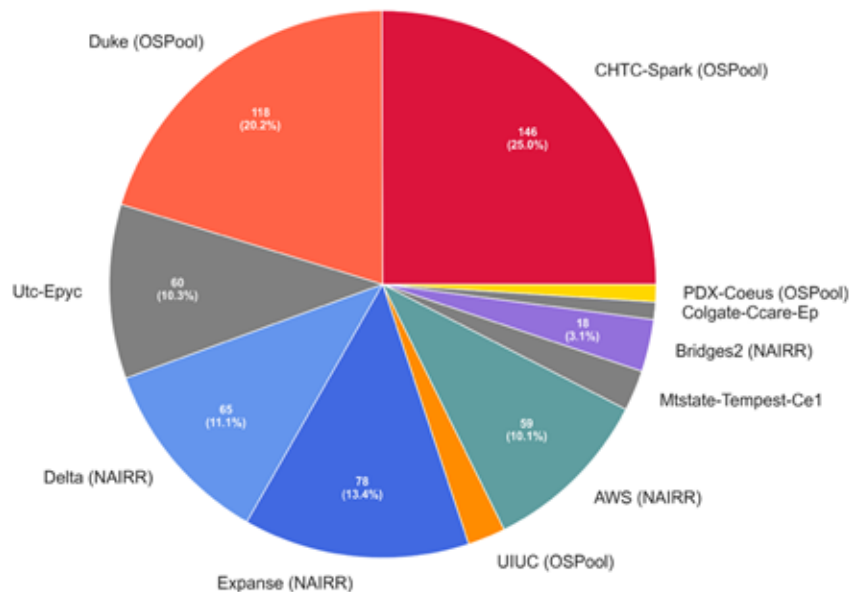
We utilized allocations at Anvil, AWS, Bridges-2, Delta, and Expanse to supplement the OSPool resources via **HTCondor Annex**, moved checkpoints and training data in via **OSDF**, and **ran epochs within NAIRR providers as if they were local**.

## Grand totals (Both model variants, heterogeneous approach)

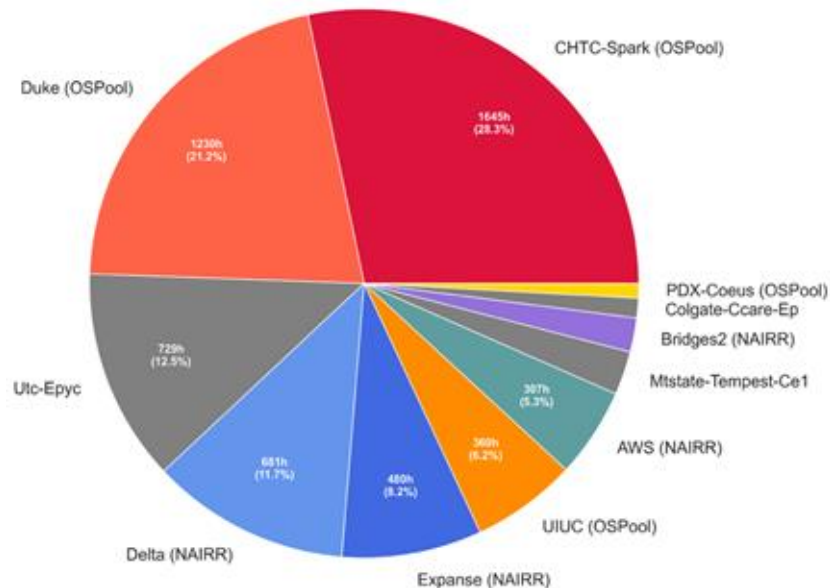
<b>Total epochs</b>	1,261
<b>Total GPU hours</b>	19,695
<b>Total data moved via OSDF</b>	~43.4 TB
<b>Unique GPU types</b>	9
<b>Unique sites</b>	6 (OSPool, Expanse, Delta, Bridges2, Anvil, AWS)
<b>Unique sites, with OSPool sites breakout</b>	13

# Per-resource contributions

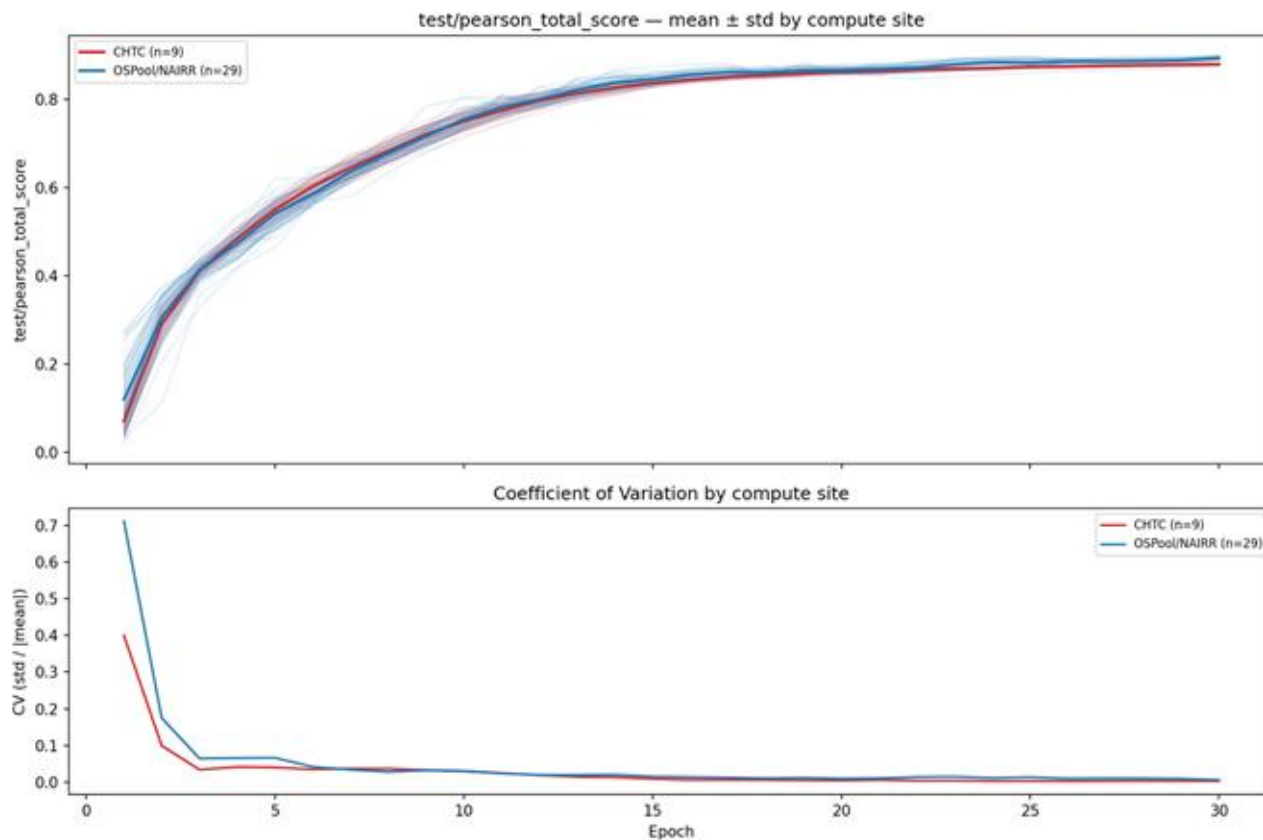
## Successful Epochs by Resource



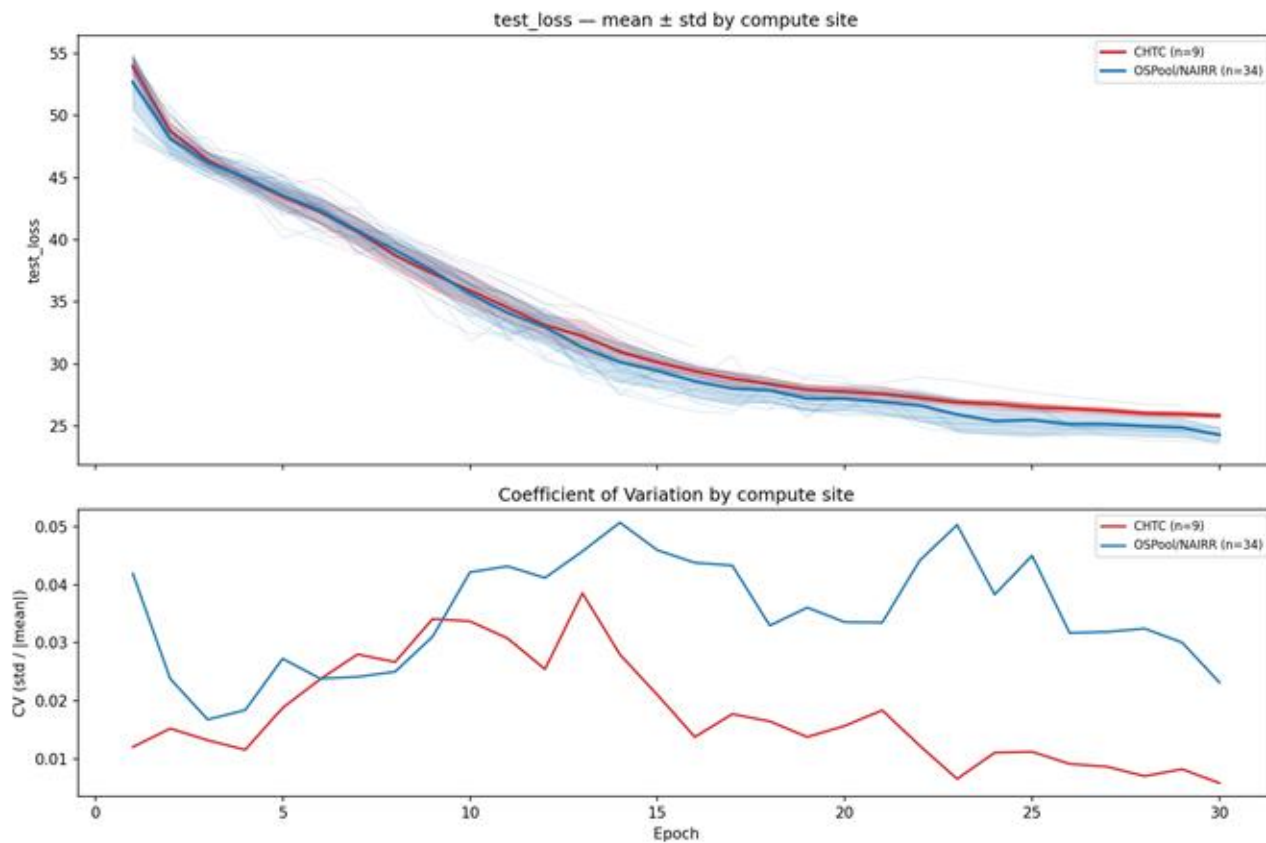
## Total Computation Time by Resource



# Observations and analysis



# Observations and analysis



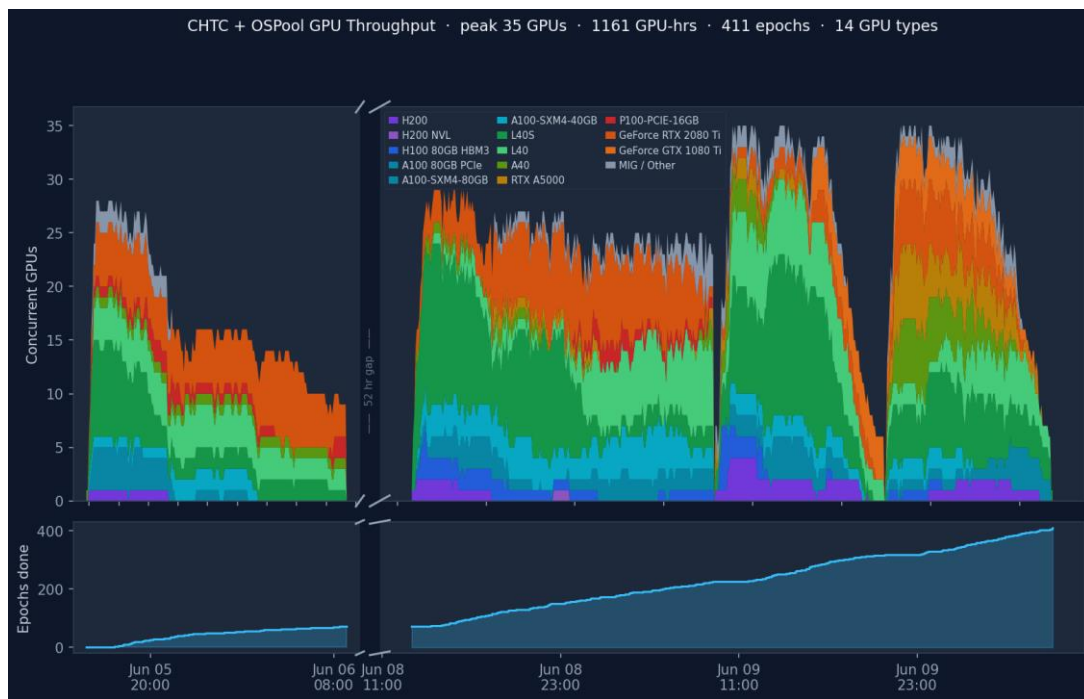
# Conclusions

- Models converge similarly in both heterogeneous and homogeneous training settings
  - Heterogeneous case has higher variance and may gain *benefit* from variability (at the cost of interpretability)
- Users can leverage nationwide compute from a single OSPool AP, *including folding NAIRR allocations in via HTCondor Annex*
  - ...but the Annex was built on top of assumptions that complicate automated and flexible utilization. Work ongoing on that front!
- Next steps
  - Automation push: Full stack (simulation → **pretrain** → finetuning) pipelines for single-protein models instead of *generalized (global)* predictive model



# Single protein model training: doing more with less

- Lower VRAM (and time) requirements: lower preemption risk, orders of magnitude more GPUs available, *can be just as scientifically interesting!*



✨ Vibe-coded visualization of early throughput tests ✨

8 different protein simulations, 5 targeted model variants each

---

# High Throughput Machine Learning

- **The PATH to NAIRR is open.**
  - HTCondor Annex + OSDF let us use allocations from single AP
  - Gaps remain and rough edges exist, but integration is proven and we know where the rough edges are
- **HTC principles hold for ML!**
  - Submit locally, run globally.
  - Epoch-sized jobs can fit almost everywhere
    - Especially with checkpointing
  - Moving the work (and data) to the compute
    - Even leveraging heterogeneous GPU in a training process
  - Always have work ready to run
- **Come talk to us! Whether you're just starting to think about ML workloads or are outgrowing OSPool, we can get you going!**