

Advancing Scientific Research with AI: Collaborative Efforts in Training Large Foundation Models on the Frontier Supercomputer

Prasanna Balaprakash

Director of AI Programs

Oak Ridge National Laboratory


ORNL is managed by UT-Battelle LLC for the US Department of Energy
Sajal Dash, Isaac Lyngus, Junqi Yin, Xiao Wang,
Romain Egele, Austin Elis, Guojing Cong,
Aristeidis Tsaris, Dan Lu, Feiyi Wang
Massimiliano Lupo Pasini, Jong Choi, Xiao Wang,
Pei Zhang, John Gounley, and Paul Laiu



ORNL has a rich history leveraging AI for science



Oak Ridge Applied Artificial Intelligence Project



Automated machines



- #1 HPL-MXP @ 10 exaflops for AI
- #2 on GREEN500
- Scaled to 1T+ parameter AI model training

1940–1970


1980

1990

2000



supports spectroscopy, environmental management, nuclear fuel reprocessing, and programming assistance



Robotics



Summit. World's "smartest" supercomputer optimized for AI

AI transforming science and national security


Accelerating scientific discovery, fortifying energy infrastructure, and enhancing national security



**Spallation
Neutron Source**



**Manufacturing
Demonstration Facility**



**Center for Structural
Molecular Biology**



**Oak Ridge Leadership
Computing Facility**



**Cyber Science
Research Facility**



**High Flux
Isotope Reactor**

Grand challenges in AI for science

nature

[Explore content](#) ▾ [About the journal](#) ▾ [Publish with us](#) ▾ [Subscribe](#)

[nature](#) > [editorials](#) > [article](#)

EDITORIAL | 27 September 2023

AI will transform science – now researchers must tame it

A new *Nature* series will explore the many ways in which artificial intelligence is changing science – for better and for worse.

OCTOBER 30, 2023

Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence

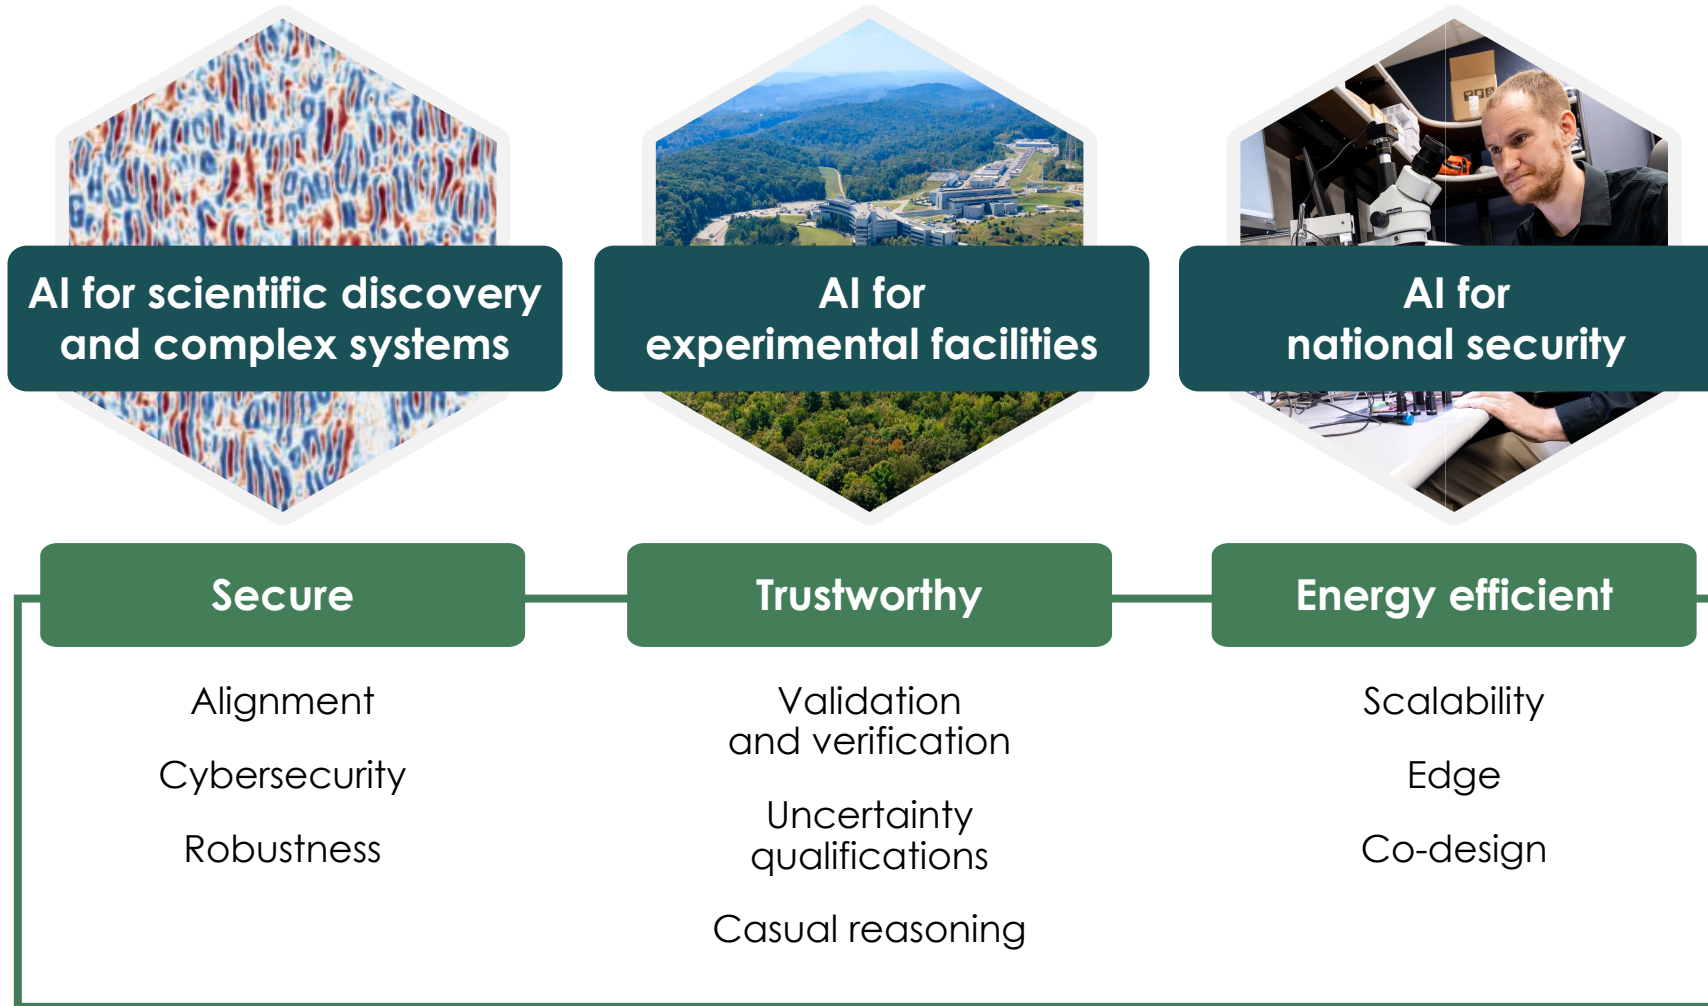
 > [BRIEFING ROOM](#) > [PRESIDENTIAL ACTIONS](#)

By the authority vested in me as President by the Constitution and the laws of the United States of America, it is hereby ordered as follows:




ORNL's AI initiative

Secure, trustworthy, and energy-efficient AI



Resources needed to train LLM-like models

One layer has $11d^2$ parameters

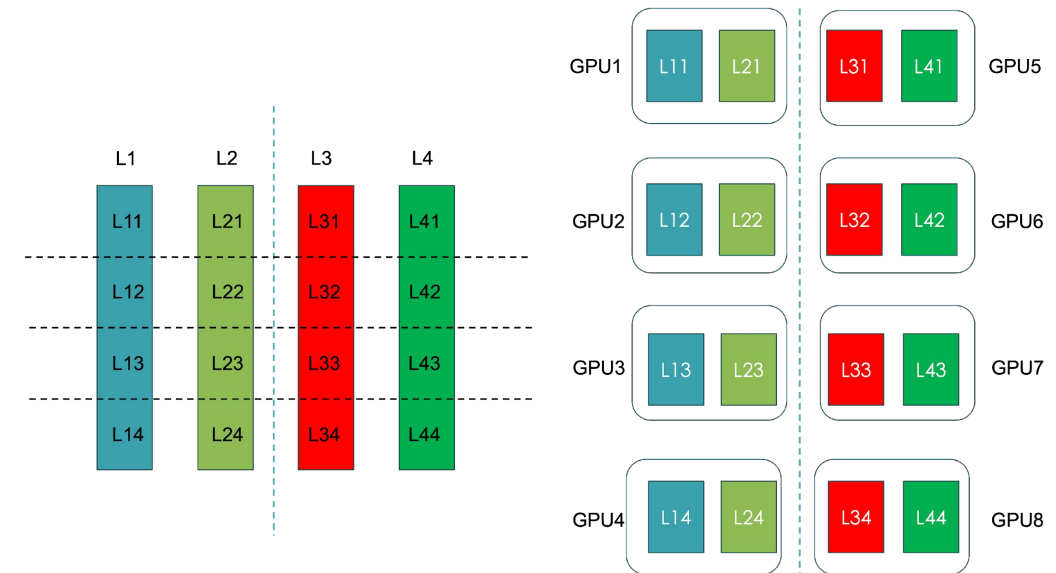


- GPT like models grows linearly with #Layers and quadratically with #Hidden dimensions
- Total parameters is roughly $12 * L * d^2$
- Training one parameter needs 20 Bytes of memory and 120 FLOPs
- Need a minimum of 8, 55, and 313 MI250X GCDs to fit models of size 22B, 175B, and 1T

Values	Memory Requirement		
	22B Model	175B Model	1T Model
Parameters (6x)	132 GB	1050 GB	6 TB
Gradients (4x)	88 GB	700 GB	4 TB
Optimizer States (8x)	176 GB	1.4 TB	8 TB
Total Memory (20x*)	440 GB	3.5 TB	20 TB



Optimizing distributed training strategies

- Ported SOTA LLM training frameworks and libraries to Frontier (Megatron-DeepSpeed, FlashAttention 1&2)
- GPT model can be trained with a combination of Tensor (TP), Pipeline (PP), and Data parallelism (DP)
- TP slices model horizontally, PP vertically, and DP replicates
- TP should limit within node, PP should use large #micro-batches

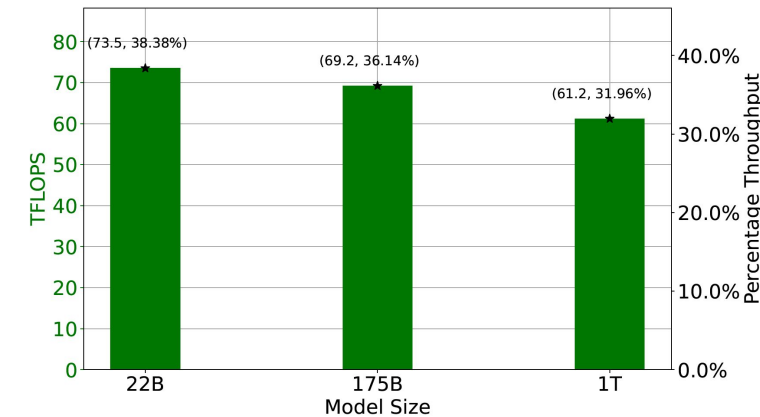


Training LLMs with 1 trillion parameters

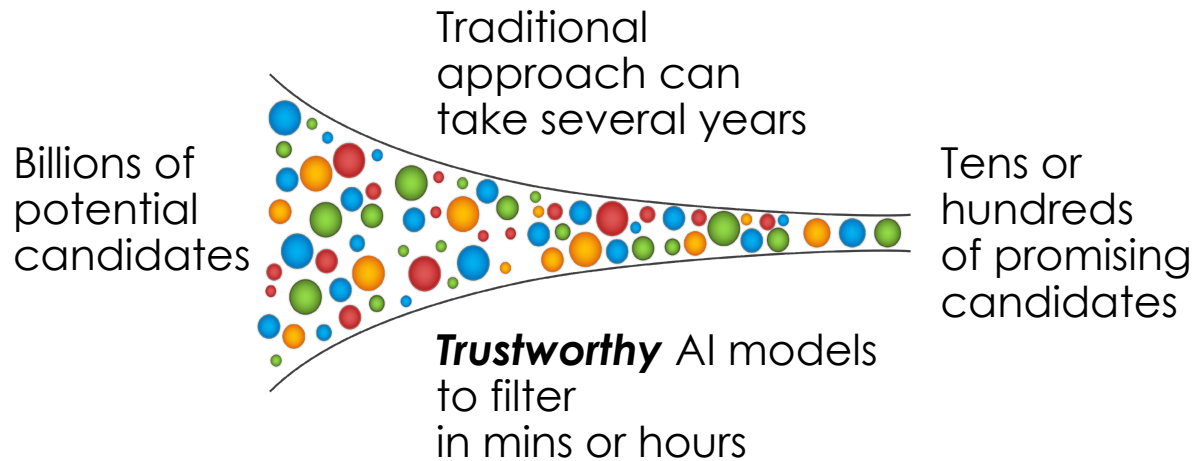
- Automatically identify the best distribution strategies to achieve high GPU throughput and scaling efficiency
- Achieved 100% weak scaling efficiency and 87.05% strong scaling efficiency at 3072 GPUs.
- Trained the models for 10-100 iterations instead of training till completion.




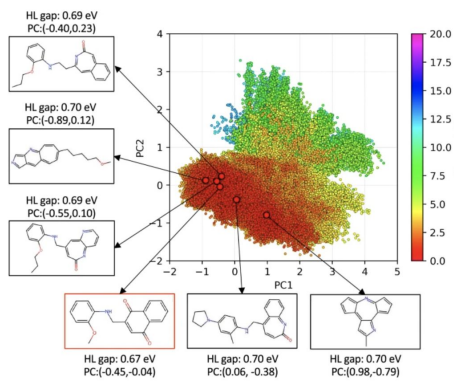
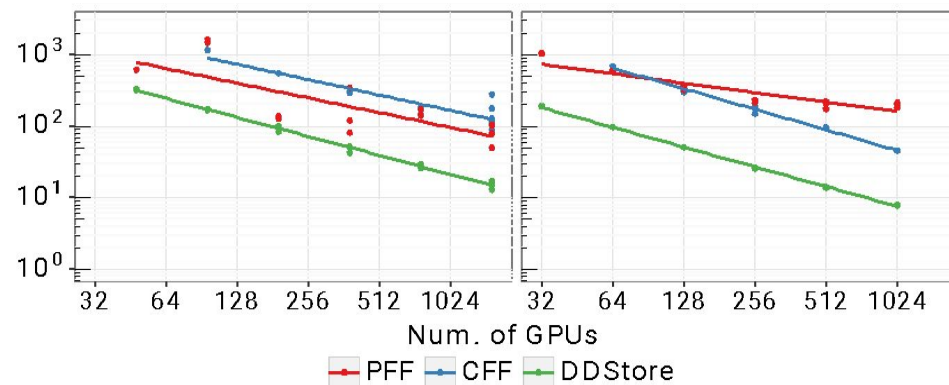
Hyperparameters	Value	
	175B Model	1T Model
TP	4	8
PP	16	64
MBS	1	1
GBS	640	1600
ZeRO Stage	1	1
Flash Attention	v2	v2
Precision	fp16	fp16
checkpoint-activations	True	True




Accelerated materials discovery via trustworthy AI models on Frontier



AI model training on 32 GPUs takes 100 hours
 AI model training on 1,024 GPUs takes 10 hours

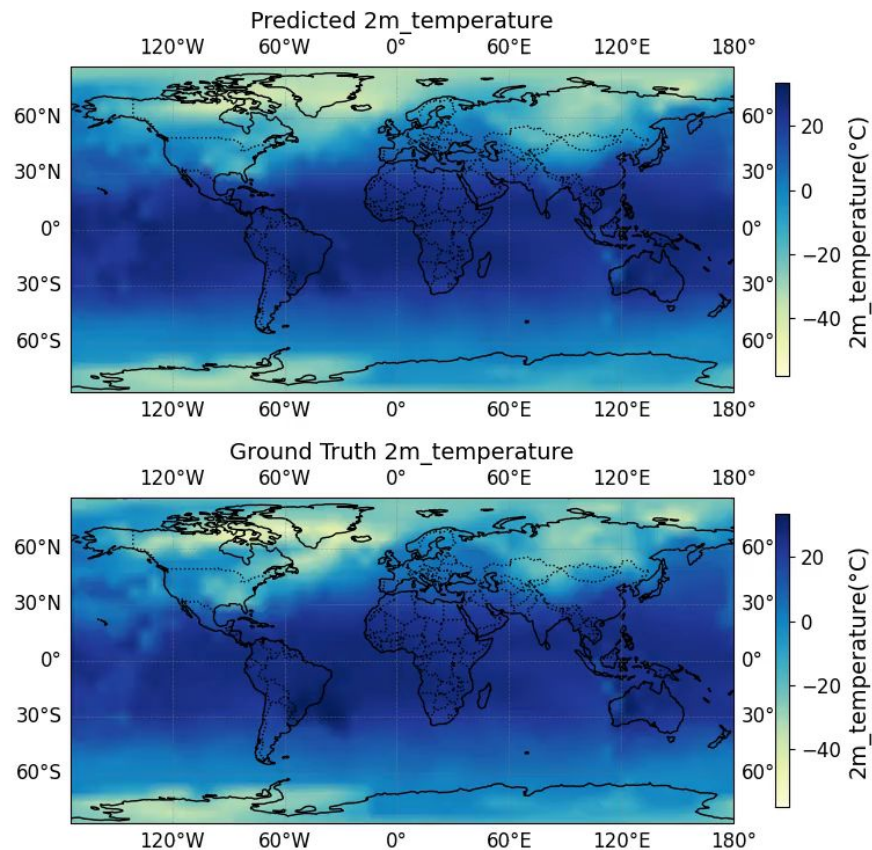


High-entropy alloy
Aerospace, mechanical manufacturing, biomedicine




Foundation model for climate and weather forecasting



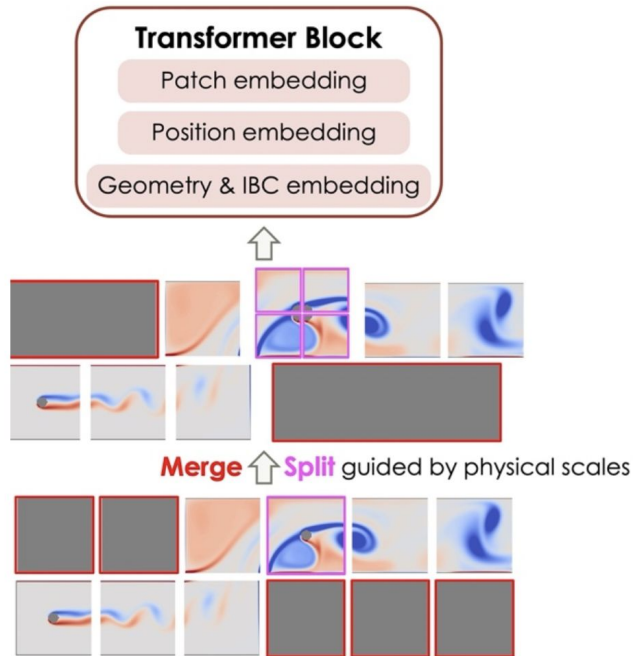
Variable 2m_temperature, at time: 2017-01-04 02:00, lead time: 72 hrs



- Scaled our AI foundation model with 1B to 5B parameters on Frontier supercomputer;
- 87% scaling efficiency.
- Our AI foundation model accurately forecasts weather 72 hours ahead.



Foundation models for fluid dynamics



- Multiscale vision transformer models □ efficient representation & computing resources distribution
 - Adaptivity guided by physics, like Adaptive Mesh Refinement (AMR)
 - Hierarchical transformers: Transformers at various scales with a two-way coupling

Data credits:
1. Miniweather by Matt Norman (ORNL),
<https://github.com/mnorman/miniWeather>
"CGL @ ETHZ - Data."
2. <https://cgl.ethz.ch/research/visualization/data.php>
(accessed Jun. 08, 2023).

Frontier first indicates path for energy efficient scaling

- First time in open science setting and on non-NVIDIA hardware
 - So far, large runs on only NVIDIA hardware
- No heroic effort in software engineering
 - Largely dependent on existing software like Megatron-DeepSpeed
- Train massive LLMs with good efficiency using an entirely NVIDIA-free supercomputer
- Don't need an army of researchers at a national lab to make productive use of AMD GPUs for training large LLMs

Frontier trained a ChatGPT-sized large language model with only 3,000 of its 37,888 Radeon GPUs — the world's fastest supercomputer blasts through one trillion parameter model with only 8 percent of its MI250X GPUs

News By Matthew Connatser published January 07, 2024

Now you're playing with AI power!

Comments (19)



12 OAK RIDGE National Laboratory | COMPUTING FACILITY

(Image credit: ORNL)

Researchers at Oak Ridge National Laboratory trained a large language model (LLM) the size of ChatGPT on the Frontier supercomputer and only needed 3,072 of its 37,888 GPUs to do it. The team published a research paper that details