



Does scientific computing have anything to offer scientific foundation models?

Michael W. Mahoney

ICSI, LBNL, & Dept of Statistics, UC Berkeley

May 2025



Physics Informed Learning?

Computational Science is an important tool that we can use to incorporate physical invariances into learning, but until recently it was missing from mainstream ML.

“**Computational Science** can **analyze past events** and **look into the future**. It can explore the effects of thousands of scenarios for or in lieu of actual experiment and be used to study events beyond the reach of expanding the boundaries of experimental science”

—Tinsley Oden, 2013

To make further progress in ML it is crucial that we incorporate computational science into learning.



J. Tinsley Oden's Commemorative Speech: "THE THIRD PILLAR: The Computational Revolution of Science and Engineering", Honda Prize, 2013.

Who has digested “the bitter lesson” of AI/ML?

"general methods that leverage computation are ultimately the most effective, and by a large margin."

"Most AI research has been conducted as if the computation available to the agent were constant (in which case leveraging human knowledge would be one of the only ways to improve performance)"

"Seeking an improvement that makes a difference in the shorter term, researchers seek to leverage their human knowledge of the domain"

"only thing that matters in the long run is the leveraging of computation."

"These two need not run counter to each other, but in practice they tend to.

- Time spent on one is time not spent on the other.
- There are psychological commitments to investment in one approach or the other.
- And the human-knowledge approach tends to complicate methods in ways that make them less suited to taking advantage of general methods leveraging computation."

“building in how we think we think does not work in the long run ...”

E.g., computer chess, computer Go, speech recognition, computer vision, time-series forecasting, ... science!

What is a foundation model?

General-purpose technologies that can support a diverse range of use cases.

Built using **well-established techniques** from ML:

- NNs, self-supervised learning, transfer learning, etc.

New paradigm in ML:

- general-purpose models are “**reusable infrastructure**,” instead of bespoke/one-off solutions
- **building** foundation models is highly resource-intensive (100M - 1B USD, people, data, compute)
- **adapting** a foundation model for a specific use case or using it directly is much less expensive.


Term was created/popularized by Stanford Institute for Human-Centered Artificial Intelligence (HAI) Center for Research on Foundation Models (CRFM):

- Bommasani et al. “On the Opportunities and Risks of Foundation Models” arXiv:2108.07258.

How to view Scientific Machine Learning

Vertical vs Horizontal integration

If a vertical integration occurs when a company acquires a company or asset at a different part of the supply chain, horizontal integration occurs when a company consolidates with the acquisition of a company or asset at the same points of the supply chain.



How to view Scientific Machine Learning

ML is a "horizontal":

- Provides a standard applicable across multiple cross-areas
- Like the iphone, or roads/railroads, or energy infrastructure, or HPC

Domain Sciences are "verticals":

- They own domain acquisition, insight, analysis, interpretation, etc.
- You need to be a domain expert to push state of the art

High-profile successes of SciML have taken place in industry:

- "Horizontal" companies that provide tech platforms and have lots of ML expertise
- Not "vertical" companies that know one science domain and use ML for that one goal

What do business leaders care about?

- No CEO cares about ML; they care about money
- Winners are those who invest heavily in this "means to an end" ML infrastructure

What might be possible with a meaningfully *Scientific* FM?

Train on data from:

- **Atmosphere**: Climate and Weather processes
- **Land**: Water and Ecosystem processes
- **Subsurface**: Heterogeneous flows and seismicity
- **Language models**: e.g., if you want to learn 1/f noise

Transfer learn on data from:

- **Astronomy**: to discover habitable exoplanets
- **Materials Science**: to learn physics across scales in an end-to-end way
- **Chemistry**: to learn interatomic potentials for MD simulation
- **Fire/Floods/Etc.**: to learn distributions of extreme events well enough to create insurance markets
- **Nuclear Physics**: to learn classified data from public data

How is this even possible? My data are special/unique?

- No; NOT so.
- You are NOT so unique/special: ML algorithms predict movies you watch better than you do

Just call ChatGPT? Or apply the M.O. of ML to Science?

Option 1:

- **Ask** ChatGPT (or whatever LLM), post fine-tuning, to hypothesize new drugs, or what comes after the Top quark, or ...

Option 2:

- **Use** ChatGPT embeddings in a model for some other scientific objective.

Option 3:

- **Understand** the *methodology of ML**
- **Apply** that methodology to Scientific data
- **Multi-modal** Scientific data could be **text**
- It could be **simulation**, **experiment**, etc.
- Incorporate **spatio-temporal** inductive biases into **architecture** and **compute**
- Develop **foundations** for SciML



Can AI Foundation Models Drive Accelerated Scientific Discovery?

The M.O. of ML: Can AI Foundation Models Drive Accelerated Scientific Discovery?

NOVEMBER 10, 2023

By Carol Pott

Contact: cscoms@lbl.gov

Pre-trained artificial intelligence (AI) foundation models have generated a lot of excitement recently, most notably with Large Language Models (LLMs) such as GPT4 and ChatGPT. The term "foundation model" refers to a class of AI models that undergo extensive training with vast and diverse datasets, setting the stage for their application across a wide array of tasks. Rather than being trained for a single purpose, these models are designed to understand complex relationships within their training data. These models can adapt to various new objectives through fine-tuning with smaller, task-specific datasets. Once fine-tuned, these models can accelerate progress and discovery by rapidly analyzing complex data, making predictions, and providing valuable insights to researchers. The magic lies in scaling the model, data, and computation in just the right way.

**Scale data size, model size, and compute so none of them saturate, then transfer learn.*

The M.O. of ML: Foundation models for SciML?


** Scale data size, model size, and compute so none of them saturate, then transfer learn.*

Create and pre-train on diverse PDE systems



Vary/Sample all inputs (PDE coefficients, source functions, ...)

Include multiple differential operators, predict PDE solution

$$\nabla \cdot \mathbf{K} \nabla u + \mathbf{v} \cdot \nabla u + \dots = f$$




Inputs
 $f, \mathbf{K}, \mathbf{v}, \dots$




Outputs
 u

Foundation Models for SciML

Solve multiple systems using the same pre-trained model, outperforming training from scratch






Inputs
 $f, \mathbf{K}, \mathbf{v}, \dots$



PDE 1

PDE 2




PDE 3



Outputs
 u

"Towards Foundation Models for Scientific Machine Learning: Characterizing Scaling and Transfer Behavior," Subramanian, Harrington, Keutzer, Bhimji, Morozov, Mahoney, and Gholami, arXiv:2306.00258, NeurIPS23.

Model Size Increased Exponentially in 2018-22



Rethink the *design, training, inference, and role of data* for successful application of NNs in SciML

- Different than computational design for ML/LLMs in industry
- Different than computational design in HPC and scientific simulation

Amir Gholami, Zhewei Yao, Sehoon Kim, Michael W. Mahoney, Kurt Keutzer, [AI and Memory Wall](#), IEEE Micro, 2024.

Kim*, S., Hooper*, C., Gholami*, A., Dong, Z., Li, X., Shen, S., Mahoney, M.W. and Keutzer, K. SqueezeLLM: Dense-and-Sparse Quantization. *arXiv:2306.07629*

Combining domain-driven and data-driven models?

Characterizing possible failure modes in physics-informed neural networks

Aditi S. Krishnapriyan^{*,1,2}, Amir Gholami^{*,2},
Shandian Zhe³, Robert M. Kirby³, Michael W. Mahoney^{2,4}

¹Lawrence Berkeley National Laboratory, ²University of California, Berkeley,

³University of Utah, ⁴International Computer Science Institute
{aditik1, amirgh, mahoneymw}@berkeley.edu, {zhe, kirby}@cs.utah.edu

Abstract

Recent work in scientific machine learning has developed so-called physics-


Science | DOI:10.1145/3524015

Chris Edwards

Neural Networks Learn to Speed Up Simulations

Physics-informed machine learning is gaining attention, but suffers from training issues.

PHYSICAL SCIENTISTS AND engineering research and development (R&D) teams are embracing neural networks in attempts to accelerate their simulations. From quantum mechanics to the prediction of blood flow in the body, numerous teams have reported on speedups in simulation by swapping conventional finite-element solvers for models trained on various combinations of experimental and synthetic data.



Some ideas from scientific computing

Illustrative recent proof-of-principle directions:



- **SuperBench**: A Super-Resolution Benchmark for SciML
- **Traditional vs Modern ML UQ**: Over- vs under-parameterized models
- **ContinuousNet**: “numerical” convergence tests
- **ProbConserve**: a posteriori correction for conservation constraints
- **Weight Diagnostics**: WeightWatcher analysis and HTSR
- **Time Series**: LEM, ConvLEM, Chronos
- **NeurDE**: long-term prediction of nonlinear conservation laws*

*ML horizontal: kinetic theory formulation to use ML to couple between different scales

Foundational methods: SuperBench



1. A super-resolution benchmark for SciML
2. High-resolution fluid flow, cosmology, and weather datasets with dimensions up to 2048×2048
3. Pixel-level difference, human-level perception domain-motivated error metrics
4. Extensible framework



"SuperBench: A Super-Resolution Benchmark Dataset for Scientific Machine Learning," Ren, Erichson, Subramanian, San, Lukic, and Mahoney, arXiv:2306.14070

Foundational methods: SuperBench -> Update: FLEX

FLEX: A Backbone for Diffusion-Based Modeling of Spatio-temporal Physical Systems

N. Benjamin Erichson, Vinicius Mikuni, Dongwei Lyu, Yang Gao, Omri Azencot, Soon Hoe Lim, Michael W. Mahoney

We introduce FLEX (Flow Expert), a backbone architecture for generative modeling of spatio-temporal physical systems using diffusion models. FLEX operates in the residual space rather than on raw data, a modeling choice that we motivate theoretically, showing that it reduces the variance of the velocity field in the diffusion model, which helps stabilize training. FLEX integrates a latent Transformer into a U-Net with standard convolutional ResNet layers and incorporates a redesigned skip connection scheme. This hybrid design enables the model to capture both local spatial detail and long-range dependencies in latent space. To improve spatio-temporal conditioning, FLEX uses a task-specific encoder that processes auxiliary inputs such as coarse or past snapshots. Weak conditioning is applied to the shared encoder via skip connections to promote generalization, while strong conditioning is applied to the decoder through both skip and bottleneck features to ensure reconstruction fidelity. FLEX achieves accurate predictions for super-resolution and forecasting tasks using as few as two reverse diffusion steps. It also produces calibrated uncertainty estimates. Resolution 2D turbulence data show that FLEX outperforms strong baselines including unseen Reynolds numbers, physical observables (e.g., fluid flow

Subjects: **Machine Learning (cs.LG)**; Artificial Intelligence (cs.AI)

Cite as: [arXiv:2505.17351](https://arxiv.org/abs/2505.17351) [cs.LG]

(or [arXiv:2505.17351v1](https://arxiv.org/abs/2505.17351v1) [cs.LG] for this version)

<https://doi.org/10.48550/arXiv.2505.17351> 

"FLEX: A Backbone for Diffusion-Based Modeling of Spatio-temporal Physical Systems," Erichson, Mikuni, Lyu, Gao, Azencot, Lim, Mahoney, arXiv:2505.17351




Figure 1: FLEX is a backbone for modeling spatio-temporal physical systems using diffusion models. It learns residual corrections conditioned on task-specific inputs (e.g., low-resolution or past states) and physical

Foundational methods: Traditional vs Modern ML UQ

Traditional UQ versus Modern UQ in overparameterized vs underparameterized models




Figure 3: **Bagged random feature classifiers.** Black dashed line represents the interpolation threshold. Across all tasks, DER and EIR are maximized at this point, and then decrease thereafter.






Figure 4: **Random forest classifiers.** Black dashed line represents the interpolation threshold. Across all tasks, DER and EIR are maximized at this point, and then remain constant thereafter.



(a) Without LR decay.



(b) With LR decay.

Figure 5: **Large scale studies of deep ensembles on ResNet18/CIFAR-10.** We plot the DER and EIR across a range of hyper-parameters, for two training settings: one with learning rate decay, and one without. The black dashed line indicates the *interpolation threshold*, i.e., the curve below which individual models achieve exactly zero training error. Observe that interpolating ensembles attain distinctly lower EIR than non-interpolating ensembles, and correspondingly have low DER (< 1), compared to non-interpolating ensembles with high DER (> 1).

"The Interpolating Information Criterion for Overparameterized Models," Hodgkinson, van der Heide, Salomone, Roosta, and Mahoney, arXiv:2307.07785

"When are ensembles really effective?," Theisen, Kim, Yang, Hodgkinson, and Mahoney, arXiv:2305.12313, NeurIPS23

"Monotonicity and Double Descent in Uncertainty Estimation with Gaussian Processes," Hodgkinson, van der Heide, Roosta, and Mahoney, arXiv:2210.07612, ICML23

Foundational methods : ContinuousNet

1. Convergence test based on numerical analysis theory
2. Verifies whether a model has learned an underlying continuous dynamics
3. Good for super-resolution, iterative dynamics, etc.
4. Applies to NNs, SINDy, etc.




Figure 4: Double gyre fluid flow: Reconstructing fine-scale flow fields from coarse training




Figure 2: Illustration of our convergence test with different ODE-Nets. (a) Schematic of an ODE-Net

"Learning continuous models for continuous physics," Krishnapriyan, Queiruga, Erichson, and Mahoney, arXiv:2202.08494, Comm Phys (2023)
 "Continuous-in-Depth Neural Networks," Queiruga, Erichson, Taylor, and Mahoney, arXiv:2008.02389

One way to address failure modes: ProbConserve

1. Compute mean and variance estimates
2. Update model (with oblique projection, depending on heteroscedasticity structure)
3. Good for sharp discontinuities






Figure 1: Illustration of the “easy-to-hard” paradigm for PDEs, for the GPME family of conservation equations: (a) “easy” parabolic smooth (diffusion equation) solutions, with constant parameter $k(u) = k \equiv 1$; (b) “medium” degenerate parabolic PME solutions, with nonlinear monomial coefficient $k(u) = u^m$, with parameter $m = 3$ here; and (c) “hard” hyperbolic-like (degenerate parabolic) sharp solutions (Stefan equation) with nonlinear step-function coefficient $k(u) = \mathbf{1}_{u \geq u^*}$, where $\mathbf{1}_{\mathcal{E}}$ is an indicator function for event $\mathcal{E}$.




Foundational methods : Weight Diagnostics

Use methods from disordered systems theory, random matrix theory and statistical physics to **diagnose** practical problems in state-of-the-art neural networks

- “Predicting trends in the quality of state-of-the-art neural networks without access to training or testing data,” Martin, Peng, and Mahoney, arXiv:2002.06716 (2020)
- “Statistical Mechanics Methods for Discovering Knowledge from Modern Production Quality Neural Networks, Martin and Mahoney,” KDD (2019)
- “Traditional and Heavy-Tailed Self Regularization in Neural Network Models, Martin and Mahoney,” ICML (2019)
- “Heavy-Tailed Universality Predicts Trends in Test Accuracies for Very Large Pre-Trained Deep Neural Networks,” Martin and Mahoney, SDM (2019)
- “Implicit Self-Regularization in Deep Neural Networks: Evidence from Random Matrix Theory and Implications for Learning,” Martin and Mahoney, arXiv:1810.01075 (2018)
- (<https://github.com/CalculatedContent/ww-trends-2020>)

Analyzing DNN Weight matrices with **WeightWatcher**



Foundational methods : Time Series

Published as a conference paper at ICLR 2022

LONG EXPRESSIVE MEMORY FOR SEQUENCE MODELING

T. Konstantin Rusch
ETH Zürich
trusch@ethz.ch

Siddhartha Mishra
ETH Zürich
smishra@ethz.ch

N. Benjamin Erichson
University of Pittsburgh
erichson@pitt.edu

Michael W. Mahoney
ICSI and UC Berkeley
mmahoney@stat.berkeley.edu





Figure 1: High-level depiction of CHRONOS. (Left) The input time series is scaled and quantized to obtain a sequence



"Chronos: Learning the Language of Time Series," Ansari et al., arXiv:2403.07815

"Using Pre-trained LLMs for Multivariate Time Series Forecasting," Wolff, Yang, Torkkola, and Mahoney arXiv:2501.06386

"Long Expressive Memory for Sequence Modeling," Rusch, Mishra, Erichson, and Mahoney, arXiv:2110.04744, ICLR22

Foundational methods : NeurDE

NeurDE: Neural Discrete Equilibrium

- Applicable to nonlinear conservation laws
- Use kinetic theory to “lift” system into higher-dim
 - Swaps nonlocal nonlinearities for linear system with a single nonlinear function
 - Tracks single particle distributions by adding velocities
- Simple NNs into the nonlinear function (Fig. 2a)
 - Predicts equilibrium particle distribution at a given time and position with two five-layer MLPs
- Linear system naturally facilitates splitting temporal evolution
- NeurDE learns underlying physics and *can predict supersonic shocks* (Fig. 2)
 - We train on early times for a single trajectory (b)
 - We predict equilibrium features not present in training data (c&d)

NeurDE spans micro-to-macro scales

- It calculates microscopic densities, but ultimately predicts macroscopic flows




Fig 1: Relation between lifting, splitting, and use of NeurDE






Fig 2: Illustration of NeurDE and supersonic predictions

Foundational methods: can be useful in your vertical ...

... for science:

- Earthquake early warning: to turn off critical infrastructure
- Scientific GenAI: to uncover physically-meaningful data ground motions
- Etc.



"Learning Physics for Unveiling Hidden Earthquake Ground Motions via Conditional Generative Modeling," Ren, Nakata, Lacour, Naiman, Nakata, Song, Bi, Malik, Morozov, Azencot, Erichson, and Mahoney, arXiv:2407.15089

"WaveCastNet: An AI-enabled Wavefield Forecasting Framework for Earthquake Early Warning," Lyu, Nakata, Ren, Mahoney, Pitarka, Nakata, and Erichson, arXiv:2405.20516

SciGPT: Scalable Foundation Model for Scientific Machine Learning

Motivation: In spite of recent effort, **there is no scientific foundation model (SFM)** that:

- (1) has been trained on a broad range of data
- (2) across different domains, and space and time scales,
- (3) to gain an understanding of multiple physical processes and their interactions in a complex scientific system.

Goal: To **develop a broad-based SFM ``blueprint"** that:

- (1) is applicable via transfer learning to multiple scientific domains and
- (2) provides a clear blueprint to develop a general scientific foundation model.

Longer Term: Will provide a **clear path forward** for more general investment:

- (1) for a general scientific foundation model and
- (2) for multiple domain-specific scientific ML models.

Data Sets




Figure 2: Example of the diverse Earth science data collected at a range of spatial and temporal scales (Figure from [131]). See also Table 1.

Datasets or Collections	Type	Spatial Extent (Resolution)	Temporal Resolution	Modality	Usage
ERA-5 [22]	Climate/weather data product	Global (25 km)	hourly	Gridded timeseries	Training
Daymet, PRISM, NARR [23]	Climate/weather data product	US (1km, 4 km, 32 km)	daily, daily, 3-hours	Gridded timeseries	Training
GHCN [23]	Climate observations	Global (point)	daily	Univariate sensor timeseries	Training
GRDC[24]	River flow observations	Global (point)	daily	Univariate Sensor Time-series	Training
FLUXNET* [25]	Land-Atmosphere energy, water flux observations	Global (point)	daily	Univariate sensor Time-series	Training
MODIS [26]	Remote sensing of land surface	Global (250 m)	daily	Images	Training
HLS [27]	Remote sensing of land surface	Global (30 m)	2-3 days	Images	Training
CMIP6/ESGF* [28]	Long-term climate simulations	Global (O(100)km)	daily, monthly	Gridded Time-series	Training (1-2 models only)
SEG Open data [29]	Seismic experiments, simulation, observation	Local (O(10)m)	milliseconds	Semi-gridded Time-series	Training
EarthScope DMC [30]	Earthquakes	Global (point)		Time-Series	Training
	Seismic experiment observations	Global (O(10)m)	milliseconds	Time-series	Training
SCEC	Earthquakes simulations	Regional (O(1km))	milliseconds	Time-series	Training
BBPlatform [31]	Observations, experiments, simulations	Pore-Global	Heterogeneous	Heterogeneous	Training and Validation
ESS-DIVE* [32]	Observations, experiments, simulations				
Energy Data eXchange* [33]	Geophysical observations and simulation	O(10)-O(1)km	milliseconds-daily	Time series, Images	Training and Validation
Geothermal Data Repository* [34]	Geophysical observations	O(10)m-O(1)km	milliseconds-daily	Time-Series, Images	Training and Validation
ARM Best estimate* [35]	Climate data product	Global (point)	hourly	Univariate sensor Time-series	Validation
ILAMB* [36]	Climate, ecosystem, water observations	Global (point)	Variable-dependent	Univariate sensor Time-series	Validation

* Data generated or served by DOE

Table 1: Data available for model training and validation including experiments, simulations, and observations across a range of spatial and temporal scales.




Figure 3: Example of diverse Earth science data from Atmosphere, Land and Subsurface.

SciGPT: Scalable Foundation Model for Scientific Machine Learning, cont.

Three main challenges: that currently block the development of a SFM:

- (1) **lack of “neural scaling”** w.r.t. model/data/compute as well as spatio-temporal scaling;
- (2) **lack of control on out-of-distribution generalization**; and
- (3) **lack of broad-based multi-modal data** for training.

Approach: **Adopt the main methodology that ML researchers do:**


- (1) used to develop CV and NLP FMs,
- (2) adapting those methods as needed to the properties of scientific data.

“Scale model and data and compute so none of them saturate, then transfer learn”


Possible SciGPT applications in $X=\{\text{Earth Sciences}\}$?

Prediction of extreme events & impacts


Earthquakes




Climate impacts on watersheds, tipping points



Learning physics across scales & Earth system components (atmosphere, land, subsurface)



E.g. power law dynamics common in natural systems



Looking forward ...

Question:

- Does scientific computing have anything to offer scientific foundation models?

Looking forward ...

Question:

- Does scientific computing have anything to offer scientific foundation models?

Answer:

- Yes, but ...