

Scalable physics-informed neural networks for large-scale simulation

Supervisor: Ben Moseley

Keywords: physics-informed neural networks, domain decomposition, random feature methods, preconditioning, finite element methods, GPU computing, large-scale simulation, scientific machine learning

Background

Scientific research relies on our ability to **simulate** physical phenomena. From modelling the Earth's climate to predicting the impact of earthquakes, simulations allow us to make predictions, test hypotheses, and explore scenarios. However, the realistic simulations we want to carry out today are often large- and multi-scale, and traditional numerical methods, such as the finite element method (FEM), can require a supercomputer for each simulation.

Scientific machine learning offers new ways of tackling these challenges [1]. Physics-informed neural networks (PINNs) [2,3] are a popular approach for carrying out simulations using neural networks: unlike traditional numerical methods, they do not require complex meshes and can easily incorporate observational data. However, PINNs are **not yet competitive** with FEM for large-scale, realistic problems. They suffer from spectral bias, which slows their convergence for high-frequency solutions, and their underlying optimisation problem is nonlinear and typically becomes harder as the problem size grows, requiring many expensive gradient descent steps.

Our recent work has made steady progress towards overcoming these limitations. Finite basis PINNs (FBPINNs) [4] split the domain into many small, overlapping **subdomains**, each with its own small network. Each network only needs to learn a simple, local part of the solution, which mitigates spectral bias and allows the networks to be trained in parallel, making high-frequency and multi-scale problems more tractable. Multilevel FBPINNs [5] add **coarser levels** of subdomains, which better share information across the domain and improve convergence as the number of subdomains grows. ELM-FBPINNs [6] then fix each network's hidden layers at random so that only the output weights are learned, transforming the expensive nonlinear optimisation problem into a **linear least-squares problem** which can be solved far more efficiently using sparse solvers. Finally, we found that **preconditioning** these linear systems speeds up their solution by orders of magnitude [7]. Together, these advances have reduced the solve times and improved the accuracy of PINNs by orders of magnitude, with solve times now approaching those of FEM. In fact, our methods are now algorithmically quite similar to FEM: both involve solving a sparse linear system, and performance depends strongly on the conditioning of that system.




Figure 1: Scalable simulation with physics-informed neural networks. The domain is divided into many overlapping subdomains, and the solution is represented using localised neural networks. When using random feature subdomain networks, training reduces to solving a sparse linear system.

Project

Whilst we have made strong progress in scaling PINNs, important bottlenecks remain: achieving fast convergence with large numbers of subdomains is still challenging, efficient parallel GPU implementations of our preconditioned ELM-FBPINNs do not yet exist, and the convergence theory of our methods is not well understood.

The goal of this project is to further improve the scalability of our methods. We will design **preconditioners and solvers** for the resulting sparse linear systems, investigate how best to **initialise** the subdomain networks, and develop optimised **multi-GPU** implementations. We will also draw on ideas from FEM, such as **adaptive domain refinement**, to reduce the number of subdomains required and to improve solver convergence.

A key goal is to efficiently solve large-scale, **realistic simulations**, such as turbulent fluid simulations (for example, modelling the Earth's climate) and inhomogeneous wave simulations (for example, modelling regional earthquakes), working with domain-specialist teams across Imperial. We will also carry out a robust and fair benchmark against FEM and, given the similarities between the two approaches, **compare them in depth** to understand their core mathematical differences, with the aim of improving both fields.

Key research questions are: can we achieve scalable convergence with large numbers of subdomains? How should our sparse linear systems be preconditioned and solved, and how should the subdomain networks be initialised? How can our methods be implemented efficiently across multiple GPUs? What are their theoretical convergence properties? How do they compare with FEM in a fair benchmark, and what can each field learn from the other?

For more details: please see our [project page](#) and [GitHub repository](#).

Impact

Scalable PINNs would give scientific modellers a powerful new simulation tool, unlocking the flexibility of neural networks to solve large-scale, realistic problems. Furthermore, understanding where and why PINNs and FEM differ could yield insights that benefit both fields, helping to overcome the computational intractability of large-scale simulation in areas such as climate modelling, geophysics, and engineering.

Supervisory team

Dr. **Ben Moseley** is an Assistant Professor in AI (Schmidt AI in Science Fellow) at the Department of Earth Science and Engineering and a Fellow at the Imperial I-X Centre. He heads the Scalable Scientific Machine Learning Lab and is an expert in scientific machine learning, physics-informed neural networks, neural differential equations, hybrid modelling, learned inverse algorithms, high-performance computing, geophysics, and planetary data science.

We are open to identifying further co-supervisors or project advisors as necessary. We actively encourage collaboration with industry and other research groups.

Research group

The student will be part of the **Scalable Scientific Machine Learning Lab** headed by Dr. Ben Moseley. The lab accelerates scientific research by designing scientific machine learning algorithms and applying them to impactful problems across science. See our [lab website](#) for more information.

We are a highly **cross-disciplinary** team – we train our members across machine learning, applied mathematics, high-performance computing, and in domain-specific areas including geophysics, climate science, and planetary science. We **collaborate** with other groups at the Department of Earth Science and Engineering, I-X (Imperial's AI initiative), other Imperial departments, and with external universities and industry partners. Lab members are encouraged to present and publish at high-impact conferences and journals.

Student profile

We are looking for someone who is motivated to complete a PhD at the intersection of scientific machine learning, numerical analysis, and high-performance computing. This is an interdisciplinary project, and we do not expect candidates to arrive with expertise in all areas; instead, we are looking for someone with a strong technical foundation, enthusiasm for interdisciplinary research, and an ability to approach complex problems with creativity and curiosity.

Essential qualifications / experience:

- Good Master's degree in a relevant field (e.g. applied mathematics, physics, computer science, engineering, or related areas). Motivated candidates with an excellent bachelor's degree and a relevant research portfolio are encouraged to apply
- Completed courses in machine learning and/or applied mathematics
- Coding proficiency in e.g. Python/ C++/ Julia/ Fortran

Desirable qualifications / experience:

- Experience in numerical modelling (finite difference, finite element, spectral methods, etc)
- Experience in numerical linear algebra, iterative solvers and preconditioning, and/or domain decomposition methods
- Understanding of scientific machine learning, in particular physics-informed neural networks
- Proficiency with Python machine learning frameworks (PyTorch, JAX (with Equinox))
- Experience in scientific, HPC, GPU, and/or parallel computing
- Relevant publications and/or industry experience are a plus

Funding

This project is not currently funded through a research grant and is eligible for College and/or Departmental scholarship funding. For more details on scholarship funding and deadlines see [here](#).

Apply

If you are interested, please start by sending us:

- CV (including education, and any research experience).
- Brief motivation letter (200 – 400 words) where you should highlight how your experience enables you to pursue the project, and your own vision for the project (can be in the email body).
- Any additional materials that support your application (optional).

For more details on our lab's PhD application process see [here](#).

For more details on the Imperial PhD application process see [here](#).

Contact

Dr. Ben Moseley: b.moseley@imperial.ac.uk

References

- [1] Moseley (2022). Physics-informed machine learning: from concepts to real-world applications. PhD Thesis.
- [2] Raissi et al., (2019). Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*.
- [3] Lagaris et al., (1998). Artificial neural networks for solving ordinary and partial differential equations. *IEEE Transactions on Neural Networks*.
- [4] Moseley, et al., (2023). Finite basis physics-informed neural networks (FBPINNs): a scalable domain decomposition approach for solving differential equations. *Advances in Computational Mathematics*.
- [5] Moseley, et al., (2024). Multilevel domain decomposition-based architectures for physics-informed neural networks. *Computer Methods in Applied Mechanics and Engineering*.
- [6] Anderson, et al., (2026). ELM-FBPINNs: an efficient multilevel random feature method. *Machine Learning for Computational Science and Engineering*.
- [7] van Beek, et al., (2025). Local feature filtering for scalable and well-conditioned random feature methods. *Computer Methods in Applied Mechanics and Engineering*.