Learning fast and generalizable climate models with neural differential equations

Supervisor: Ben Moseley
Co-supervisor: Christopher Pain

Keywords: weather forecasting, climate modelling, neural differential equations, physics-informed machine learning, scientific machine learning

Background

Weather and climate simulations are becoming increasingly important as the Earth's climate changes. For example, extreme weather events are growing in frequency and magnitude, and localised simulations over hours and days are essential for mitigating their impact. Furthermore, as we formulate climate change policies, long-term global simulations are crucial for determining the impact of climate change.

However, whilst today's simulations are highly sophisticated, they require **immense computational resources**. Traditional simulations rely on building complex mathematical models of many different physical phenomena, and these must be solved numerically, requiring significant computational effort. Furthermore, given the complexity of the Earth's climate, it is likely these models **do not capture all the relevant physics** exhibited in the real world. Thus, the accuracy and number of simulations that can be run is limited.

An increasingly popular alternative is to use **machine learning** (ML)-based models for simulation. Models such as FourCastNet [1] and GraphCast [2] are orders of magnitude faster than numerical simulation and are approaching similar accuracy. By learning from data, they can also account for missing physics. However, a major downside is that they struggle to **generalise**, i.e., generate stable predictions over long times. By entirely replacing numerical models with ML models, their predictive capabilities become limited by their training data and they do not have a strong inductive bias, or **prior understanding**, of the underlying physical system.

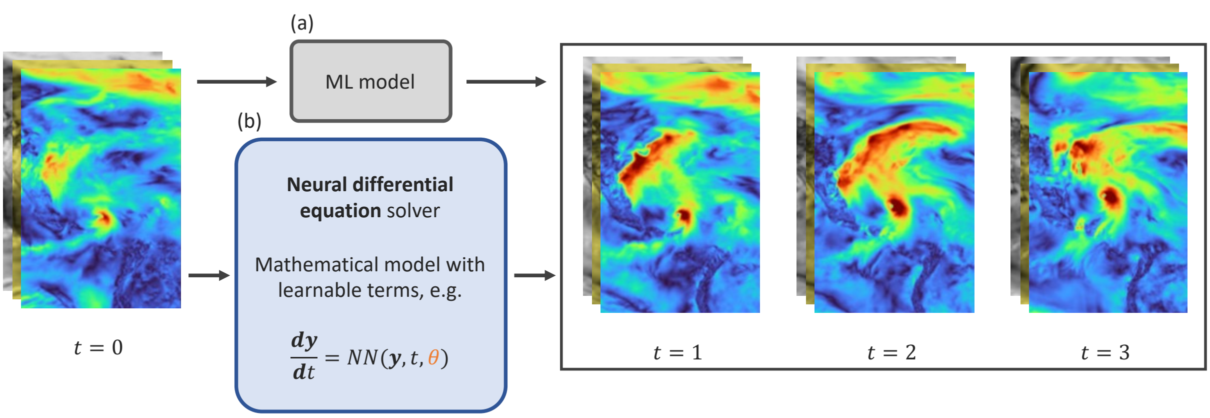

Figure 1: Workflow overview. Instead of using a purely data-driven ML model to simulate the weather and climate as shown by (a), in this project we will investigate using neural differential equations (NDEs) (b). The NDE defines a hybrid model which combines an existing mathematical model with learnable components.

Project

The goal of this project is to design hybrid weather and climate simulators which tightly **combine** traditional numerical models with ML components. By combining the best of both modelling paradigms, we will design simulators which have strong physical priors, can learn from data, and are fast, generalizable, and accurate.

We will use **neural differential equations** [3,4] to define our hybrid models (Figure 1). An NDE model is defined as follows. First, a mathematical model of the system is designed in the form of a (set of) differential equation(s), where some terms in the differential equation are treated as learnable and defined using neural networks. Then, given some initial conditions of the system, the equation is discretised and solved using a numerical solver. Finally, the parameters of the networks are learned by matching the output of the solver to some desired outputs. The NDE's underlying equations impose **physical constraints** on its outputs, whilst its learnable components allow **flexibility** to fit its training data.

A key aspect of the project will be in **investigating different NDE designs**. Recent work [5] designed an NDE for climate simulation which is built using a sophisticated general circulation mathematical model. In this project we will investigate trade-offs between the complexity of the model, the number of learnable components, its accuracy, its stability, and the time it takes to solve the model. Key research questions are, do NDEs produce more stable predictions than purely ML-based models? How fast are they compared to traditional climate models? What effect do physical constraints have on model performance?

Impact

Faster and more accurate weather and climate simulations will have a transformative impact on climate science and policy. For example, sensitivity studies of extreme weather events could be carried out in less time, allowing us to better quantify our uncertainty in their impact. Furthermore, more accurate simulations would increase our confidence when making longer-term predictions, supporting more robust policy.

Supervision

Dr. **Ben Moseley** is a Lecturer 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.

Prof. **Christopher Pain** is a Professor at the Department of Earth Science and Engineering and head of the Applied Computation and Modelling Group (AMCG). AMCG specialises in the development and application of innovative and world leading modelling techniques for earth, engineering and biomedical sciences. The group has core research interests in numerical methods for ocean, atmosphere and climate systems, engineering fluids including multiphase flows, neutral particle radiation transport, coupled fluids-solids modelling with discrete element methods, turbulence modelling, inversion methods, imaging, and impact cratering.

We are open to identifying further co-supervisors or 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 Earth science, space science, and other domains.

We value **interdisciplinarity** – 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 in scientific machine learning and climate modelling. Given the interdisciplinary nature of this project and group, we do not expect candidates with deep experience in all areas; instead, we are looking for someone who has a strong foundation, a willingness to work across disciplines, a passion for continuous learning, and the ability to approach complex problems with creativity and curiosity.

Essential qualifications / experience:

- Good Master's degree in a relevant field, e.g. mathematics, physics, computer science, engineering, or Earth science. Motivated candidates with an excellent bachelor's degree and a relevant research portfolio are encouraged to apply
- Completed courses and/or research experience in two or more of the following categories: machine learning, climate science / physics, applied mathematics / numerical modelling
- Coding proficiency in e.g. Python/ C++/ Julia/ Fortran

Desirable qualifications / experience:
- Experience in climate modelling and weather forecasting
- Understanding of scientific machine learning, neural differential equations and hybrid workflows
- Familiarity with different deep learning architectures
- Proficiency with Python machine learning frameworks (PyTorch, JAX (with Diffrax/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: https://www.imperial.ac.uk/earth-science/prosp-students/phd-opportunities/funding/

Apply

If you are interested, please start by sending us a CV and detailed description (~200-400 words) of your relevant experience and specific research interests.

For more details on the Imperial PhD application process see here: https://www.imperial.ac.uk/earth-science/prosp-students/phd-opportunities/

Contact

Ben Moseley will be joining Imperial in November 2024. Before this date, please visit his personal website for contact details: https://benmoseley.blog/

References

[1] Pathak et al., (2022). FourCastNet: A Global Data-driven High-resolution Weather Model. ArXiv.
[2] Lam et al., (2023). Learning skillful medium-range global weather forecasting. Science.
[3] Chen et al., (2018). Neural Ordinary Differential Equations. NeurIPS.
[4] Kidger (2022). On Neural Differential Equations. PhD Thesis.
[5] Kochkov et al., (2024) Neural general circulation models for weather and climate. Nature.