CONTINENTS

Net Zero and HPC: Leveraging Machine Learning for Atmospheric Modelling

Kara Moraw [1], Ioana Colfescu [2, 3]
[1]EPCC, University of Edinburgh, [2] National Centre for Atmospheric Science (NCAS), [3] University of St Andrews
Email: k.moraw@epcc.ed.ac.uk

Benefits of ML for Atmospheric Sciences

Machine Learning (ML) methods are ideally suited for processing large datasets, extracting intricate patterns and approximating complex processes. As such, they have been transforming weather and climate science, increasingly complementing traditional methodologies for forecasting and analysing complex atmospheric phenomena.
There is evidence that **ML-based solutions can be computationally more efficient than traditional methods while achieving comparable accuracy.** Minimising the energy that is consumed by data centres and High-Performance Computing (HPC) resources is key in the face of climate change, and leveraging Machine Learning for scientific computing applications has the **potential to significantly reduce the resource requirements of weather and climate modelling, as well as offer faster solutions to mitigate and adapt to climate change.**

Applications [1]:

Observations
Autoencoders and Generative Models used for
- Super-resolution
- Spatial data assimilation
- Temporal interpolation

Prediction and Forecasting
Convolutional NNs, Transformers and Graph NNs used for
- Weather forecasting
- Climate prediction
- Interannual variability
- Hybrid modelling

Sub-Grid Scale Parameterisation
Deep NNs, Autoencoders and Random Forests used for
- Parameter optimisation
- Equation discovery
- Emulation

https://continents-project.github.io/

Challenges of ML for Atmospheric Sciences

Building Trust
Machine Learning models in the climate domain will need to reliably predict trends over long time spans, possibly working with out-of-distribution inputs as the climate changes.
Techniques such as explainable AI (XAI) and interpretable models help build understanding for the inner workings of an ML solution. Integrating physical constraints into the architecture and training improves confidence in the model and its performance for rare events.

Environmental Cost
The energy use of ML methods varies greatly throughout the model lifecycle. Various considerations along the way can help reduce the cost of each stage; such as suitable data hosting locations, preprocessing steps, transfer learning, and effective use of available hardware.

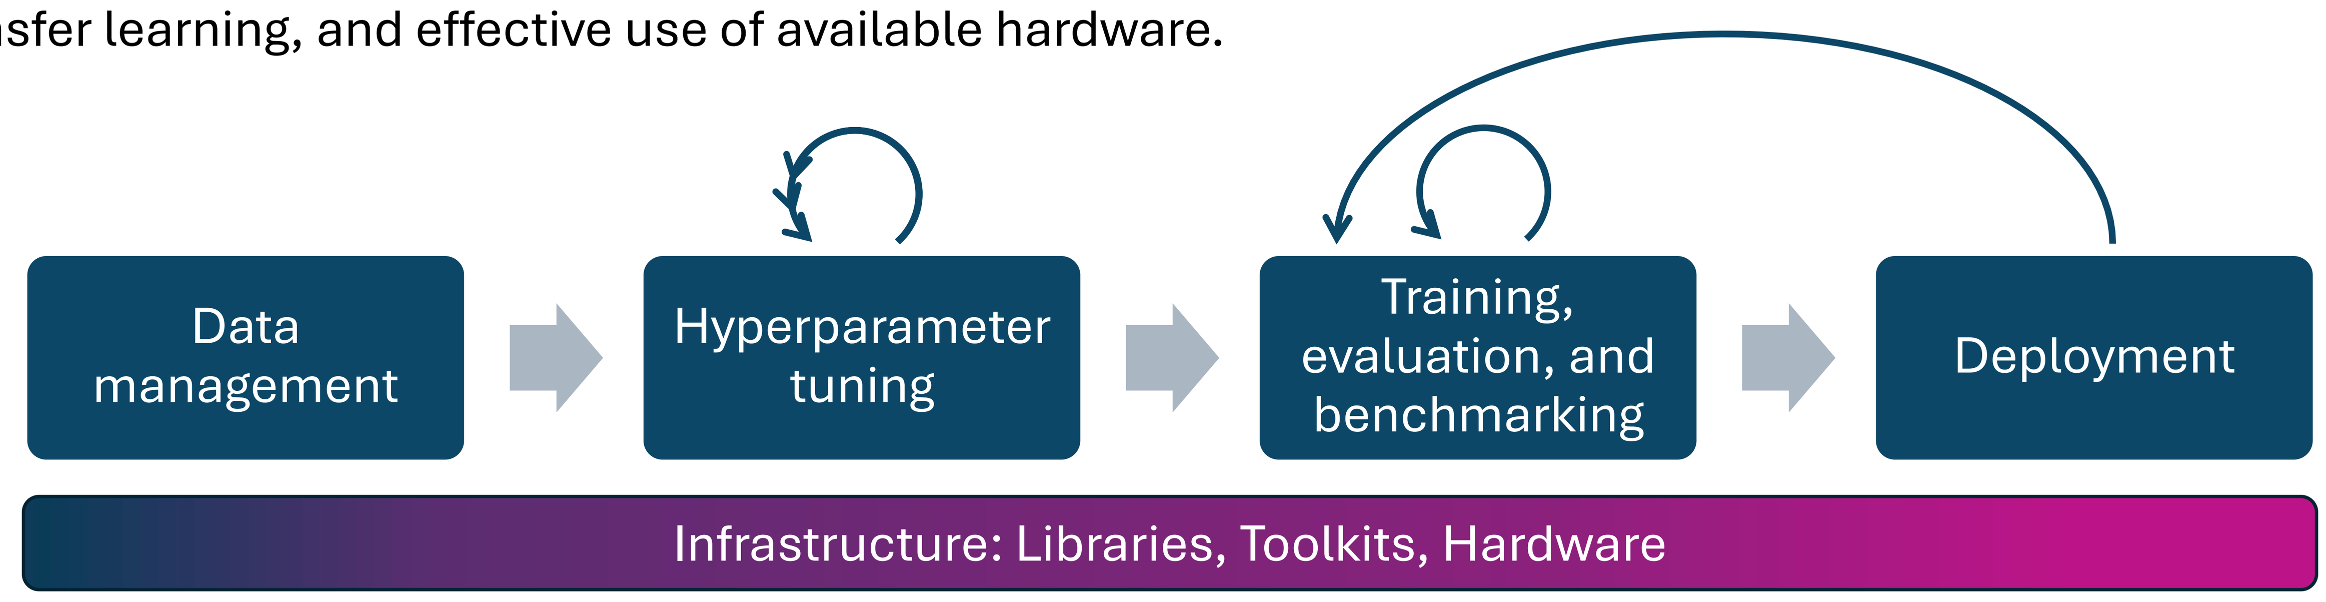

Overcoming these challenges requires interdisciplinary collaboration and knowledge exchange between domain scientists and research software engineers (RSEs).

Lifecycle adapted from [2]

Workstream 6

The CONTINENTS project is a 4 year-long collaborative programme of research. The project is led by EPCC, the UK's National Supercomputing Centre at the University of Edinburgh, in collaboration with NCAS (UK) and NCAR (USA). Workstream 6 investigates leveraging ML for the atmospheric sciences.

Objectives

- Develop methodologies and techniques for **performance-, power- and energy-efficient software** that can be deployed on a wide range of hardware. This includes exploiting ML to accelerate bottlenecks.
- **Make data a first-class citizen** of computational modelling and simulation to minimise the time and resources that are spent moving, processing, analysing and storing data. This includes preparing data to be used in ML model training.

Phases

1. Balancing the **cost** of ML training with the **benefits** of ML inference in atmospheric science: We survey the environmental impact of traditional numerical modelling and potential ML-based alternatives across various use cases. This includes the investigation of best practices and requirements both specific to the weather and climate domain, and more broadly selecting efficient ML architectures and workflows.
2. **Quantifying** the sustainability of ML in HPC and data analysis workflows: Considering the integration of ML-based components into existing large-scale modelling workflows, we will select a relevant subset of models and analyse their energy use.
3. Performance and energy **optimised use of ML** in atmospheric science workflows: We will combine the results from the quantification study with the best practices to optimise the use of ML in atmospheric science workflows.

Acknowledgements

This work is supported by the UK Research Councils through grant EP/Z531170/1.

References

[1] Bracco, A., Brajard, J., Dijkstra, H. A., Hassanzadeh, P., Lessig, C., & Monteleoni, C. (2024). Machine learning for the physics of climate. Nature Reviews Physics, 7(1), 6–20. https://doi.org/10.1038/s42254-024-00776-3
[2] Ashmore, R., Calinescu, R., & Paterson, C. (2021). Assuring the Machine Learning Lifecycle. ACM Computing Surveys, 54(5), 1–39. https://doi.org/10.1145/3453444
[logo] Earth photo by the NASA EPIC Team
[icon] Flaticon.com