Bias correcting marine ecology forecasts – from daily to seasonal scales

Ocean ecology and biogeochemistry are fundamental components of the Earth system, accounting for half the planet's primary production, playing a crucial role in global carbon cycling, and providing livelihoods for 100s of millions of people. Ocean biogeochemistry influences the wider Earth system by helping regulate ocean heat and carbon uptake and ocean acidification, and contributing to aerosols and the formation of cloud condensation nuclei. Predictions of ocean biogeochemistry and ecology are therefore vital, from daily to seasonal timescales, yet operational prediction systems remain relatively immature. This is due in part to known ecosystem model biases and a sparsity of observations, posing a challenge to data assimilation and model forecasts.

This poses the research question: "How best to combine biased and uncertain models with sparse and uncertain observations to maximise predictive skill across a range of scales, including inter-related Earth system components?" This project aims to tackle this question through combining traditional data assimilation and modelling techniques with novel machine learning methods. It will build on existing collaborations, developing the Met Office's operational prediction and reanalysis systems for ocean physics and biogeochemistry, involving the data assimilation expertise of the University of Reading, the operational forecasting and climate research expertise of the Met Office, and the ocean ecology expertise of Plymouth Marine Laboratory. The focus will be on ocean biogeochemistry, but the methods developed will be applicable to other Earth system components.

Monsoons at the kilometre scale

Monsoons supply most of the water to over two-thirds of the world's population, yet large biases in simulating monsoon rainbands have persisted over generations of models. Biases develop over only a few days, suggesting that "parameterizations" (statistical representations of sub-grid-scale processes such as tropical convection) are to blame. New models at convection-resolving scales offer the opportunity to better represent the diurnal cycle of convective rainfall, its coupling with the winds and interaction with mountains and coastlines. These models include kilometre-scale infrastructure being developed at the Met Office and ECMWF, which can more faithfully capture land-sea and mountain-valley circulations, the timing of the diurnal cycle, and storm initiation in response to subtle gradients in soil moisture.

This project will investigate errors in monsoon circulations, their scale interactions with modes of climate variability, and how these change when modelled at the kilometre scale. At first, the student will analyse models at a range of resolutions, to understand whether monsoon behaviour improves when represented at the kilometre scale. Next, the student will diagnose how tropical variability, such as the Madden-Julian Oscillation, alters monsoon circulations and whether the nature of these interactions changes at higher resolution. Finally, experiments will be designed to test the sensitivity to changes in the representation of mountains and other aspects of the surface. This will help identify the crucial processes for simulation of monsoons and their feedbacks onto the global circulation.

Using Machine Learning to obtain novel parameter sets for innovative theories of soil water- and heat transfer

Accurately predicting the weather requires an understanding of how the land surface and atmosphere interact. These interactions occur via the topsoil and vegetation exchanging energy and water with the atmosphere. Soil hydro-thermal properties, such as a soil's capacity to retain water and its ability to conduct heat, play a significant role in these exchanges, yet current weather prediction models rely on inadequate equations and parameter sets that need urgent improvement.

You will collaborate with leading scientists to develop innovative methods, using physics-informed machine learning (PIML), to create global, high-resolution soil hydro-thermal property datasets, fit for next-generation hydro-thermal theories. Normally, these properties are derived from small soil samples, using laboratory techniques. Here, in a novel approach, you will use time-series of in-situ data (e.g. soil moisture and temperature), as well as satellite imagery, to constrain standard soil water and/or heat transport equations via PIML, to derive soil properties at plot to field scales, more commensurate with km-scale weather forecasting. You will harmonise these lab-to-field scale hydro-thermal data, combine them with global environmental co-variates, and use data-driven ML techniques to create global parameter maps describing the shape of the hydro-thermal curves as a function of soil saturation. These new parameter data sets will be critical for improving weather prediction models; you will evaluate them via simulations with the ecLand model.

Optimising hybrid synthetic and real data driven predictive modelling of UK Windstorms

UK windstorms pose high-impact risks for critical infrastructure, insurance portfolios, and civil protection. However extreme instances of windstorms are under-represented in the available reference dataset ERA5 and so synthetic models are required to balance the dataset to enable the training of best performant predictive models of windstorms.

The research hypothesis is that a carefully curated hybrid dataset that fuses real observations with a task-optimised mix of synthetic outputs from WGANS+PS, Diffusion, and NeuroGCM will improve extreme-event detection, severity estimation, and spatial footprint prediction, while ensuring physics-informed interpretability and reliable uncertainty for operational use.

Thus, essentially this study comprises optimisation at two levels:

1) To develop the most effective hybrid synthetic dataset as an optimal mix of synthetic output from best performing synthetic data engines – as supplementary to the ERA5 dataset.
2) To use the above dataset to train and benchmark predictive models so as to establish the most optimal windstorms predictive modelling pipeline with physics-informed interpretability

Accordingly, this PhD study is to benchmark synthetic data quality from WGANS+PS, Diffusion Models, and NeuroGCM and to optimise their hybrid mixing with real data for best performing predictive models of UK windstorms.

The PhD study will thus develop and validate:

- A synthetic benchmarking framework for physics consistency, extreme-tail fidelity and overall predictive modelling pipeline performance enhancement.
- A hybrid data optimisation (mixing ratios optimisation) selecting the most downstream-effective combination of synthetic outputs from the best-performing engines.
- A physics-informed, interpretable ML pipeline with calibrated uncertainty and robust performance for stakeholder risk assessment decision support.

High-Resolution Modelling of Climate-Change Impacts on the West African Monsoon

Severe climate-related hazards, including droughts, floods, and episodes of heavy rainfall, are increasingly affecting West Africa. These events are becoming more frequent and intense as climate change progresses. Consequently, variations in monsoon rainfall have wide-reaching impacts on millions of people. Yet future changes in monsoon system remain highly uncertain, as climate models produce a wide range of projections for the end of the 21st century.

The sources of these uncertainties are not fully understood, but many studies point to persistent challenges in simulating precipitation and atmospheric circulation. Low-resolution models rely on parameterised convection, which limits their ability to capture the processes that drive monsoon rainfall. Increasing model resolution allows deep convection to be resolved explicitly, offering a more realistic simulation of precipitation.

How does horizontal resolution impact West African monsoon simulation and projection?

In this project, the student will use a high-resolution 2.6-km simulation developed by the Centre National de Recherche Météorologiques (CNRM) as part of the CMIP7 fast-track efforts. This simulation spans, for the first time, both historical and future periods up to 2100, enabling a detailed assessment of climate-change impacts on monsoon precipitation. Results will be compared with lower-resolution simulations (12km, 25 km and 50 km) with the same model and with many observations, reanalyses and CMIP models. Particular attention will be given to synoptic and mesoscale activity, precipitation coherence, intraseasonal monsoon variability, and land–atmosphere feedbacks. Complementary convection-permitting simulations from the University of Reading will further enhance understanding of future monsoon changes and support improvements in climate modelling.

Simultaneous evaluation of clouds and radiation in high-resolution models using EarthCARE

Clouds are a major source of uncertainty in both weather and climate prediction via their interaction with solar and thermal-infrared radiation. Launched in May 2024, the EarthCARE satellite offers the most detailed ever estimates of the properties of clouds and precipitation, including how much water they contain, mean particle size, whether they are composed of liquid or ice (or a mixture of both) and for the first time even the density and fall speed of snowflakes. Simultaneously, EarthCARE measures the reflected sunlight and emitted thermal radiation emerging from cloud top, which offers the exciting new opportunity to evaluate the properties of clouds in weather and climate models at the same time as their impact on radiation. This way radiation errors can be traced more directly to their cause than has been possible in the past. In this project, the student will use EarthCARE to evaluate the ECMWF model used for global weather forecasts and the Met Office model used for both high-resolution weather forecasts and global climate projections.

When problems are identified in the way clouds are represented in the models, it will be possible to modify them and perform new simulations to test not only whether they agree better with EarthCARE measurements, but also the impact on the accuracy of weather forecasts more generally.

Utilising AI emulators for subseasonal-to-seasonal forecasting: exploring sources of skill and hybrid forecasting approaches

AI emulators have led to a revolution in NWP over recent year, demonstrating skill and accuracy that matches and even surpasses the state-of-the-art dynamical models. Recent developments have shown that AI emulators can be stable at longer roll-out times presenting exciting opportunities for S2S forecasting on longer timescales. Despite recent evidence of potentially useful skill on S2S timescales (e.g. Kent et al., 2025), the sources of skill in these emulators are not clear. On monthly timescales and longer the interaction of the atmosphere with slowly varying of the climate system (i.e. the dynamical ocean and in the stratosphere) become increasingly important. Nonetheless, dynamical models demonstrate substantial limitations on these timescales (e.g. Garfinkel et al., 2024).

In this project we will focus on the dynamics underpinning predictable (and potentially predictable) signals in the extratropics and how they are represented in both initialised dynamical predictions and AI-trained emulators to interrogate the physical mechanisms underpinning the signals and their representation in the different modelling systems. We will utilise both types of system to identify common and divergent predictable signals, with a particular focus on extratropical teleconnections and jet/weather regimes. As the project continues, we will supplement this approach with parallel sensitivity experiments using the IFS (or UM if possible) and AI emulators.

Evaluating oceanic sources of predictability in subseasonal-to-seasonal forecasts

There are substantial predictability gaps from weeks 3-4 onwards in subseasonal-to-seasonal forecasts, after the predictability arising from atmospheric initial conditions has substantially reduced. Oceanic anomalies are the most important source of potential skill on these timescales. After initialisation, however, coupled forecast models tend to show distinct and pervasive SST drifts – or growing biases relative to observations – particularly notable in the tropical oceans; it is not clear how these drifts impact the associated forecasts. In this project, we will revisit the contributions of tropical SST signals and associated drifts on sub-seasonal forecasts. We will also aim to identify and quantify specific avenues of potential improvement.

Our specific focus will be on examining how the tropical SSTs (and associated drifts) affect teleconnection patterns to the extratropics on sub-seasonal timescales, including how these affect the influence of the MJO and ENSO. Recent work has shown that there is substantial predictability from ENSO that varies across the season but exhibits signal-to-noise errors also prevalent in longer-range forecasts (Garfinkel et al., 2024; O'Reilly et al., 2025).

The project will involve analysing historical reforecasts from coupled sub-seasonal forecast systems. With the student, we will test specific hypotheses shaped by the analysis of the reforecasts by performing a series of dedicated new experiments. The student will develop an understanding of atmospheric/ocean dynamics, ocean-atmosphere coupling, and associated teleconnection mechanisms and develop experience of using coupled general circulation models and running forecast simulations.

Accounting for the influence of the solar wind in lightning forecast models.

Research has shown a statistically significant influence on lightning strike rates over the UK in response to the passage of high-speed solar wind streams at Earth. While atmospheric convection and charge separation must be present to form Thunderclouds, the influence of the solar wind on the global electric circuit is a likely mechanism for the modulation of lightning rates. The arrival of solar wind streams can be forecast by models such as Reading's HUXt solar wind model, run operationally by the UK Met Office. Such modulation, while predictable, is yet unaccounted for in lightning flash rate diagnostics produced using the Met Office and ECMWF models. These diagnostics use a combination of thermodynamic, and cloud microphysical parameterisations computed from the model output to infer lightning strike rate.

The main research question to explore is: Can the passage of high-speed solar wind streams at Earth be used to improve the forecast skill of lightning rates?

To achieve this an analysis will be undertaken using archived forecast data over Northern Europe from both the Met Office and ECMWF to compute a suite of lightning strike rate diagnostics at various lead times. These diagnostics will be verified using Met Office lightning location system data, and for recent years, MTG-LI to identify how the lightning rate diagnostic skills are affected by high-speed solar wind events. The second part of the project would be to develop improved lightning rate diagnostics which use outputs from the HUXt model that have increased skill and at longer lead times.

Short range sea ice forecasting, using advanced models and new data

Over recent decades, the sea ice cover of the Arctic Ocean has reduced in extent, area, thickness and age. Climate models suggest the Arctic will become seasonally ice free within 20-30 years. Navigating the Arctic Ocean can considerably shorten trade routes, reducing carbon emissions and saving money. Shipping navigability has improved in recent years and this will continue, with increased interest in sea ice weather forecasts being inevitable.

This PhD project will investigate and develop a prototype for an unprecedentedly skilful sea ice forecast system, up to subseasonal timescales at various spatial resolutions (<50 km). This project will combine the latest developments in sea ice-ocean modelling such as melt ponds, form drag and floe size distribution (largely from UoR) with the latest satellite sea ice products, including CryoSat-2, SMOS, IceSat-2 and, for the first time, estimates of summer sea ice thickness. We will use a data assimilation scheme we developed based on the existing Parallel Data Assimilation Framework.

We will (i) improve the underlying sea ice-ocean model using assimilation of new, ground-truthed satellite sea ice data. This is done through examination of reanalysis increments in individual terms in the sea ice mass and momentum budgets, allowing model shortcomings to be identified and rectified; (ii) identify sources of predictive skill throughout the year but particularly in the navigable summer and shoulder seasons, building on our previous success using melt pond fraction to predict the sea ice minimum; and (iii) develop and assess a prototype sea ice forecast system.

Influence of kilometre-scale atmospheric variability on the oceanic submesoscale eddy field

Understanding ocean dynamics is critical for predicting climate change. Among these dynamics, submesoscale processes—small eddies 1–10 km wide living in the upper ocean layer —represent a new frontier in ocean modelling.

These eddies are known to play a critical role in controlling the depth of the surface mixed layer, regulating heat and CO2 uptake in the ocean interior, and bringing nutrients back into the surface layer, with long-term effects on climate.

Submesoscale eddies are generated by a family of instabilities that feed on the energy stored in larger-scale fronts and filaments (10-100 km). Previous studies have ignored the role of atmospheric variability in enhancing or dissipating the submesoscale field, in part because only the synoptic atmospheric variability, which is much larger than 1-10 km, has been considered.

As coupled ocean-atmosphere models move toward higher resolutions, the scales of oceanic submesoscale eddies and that of the newly resolved (kilometre-scale) atmospheric variability will converge. Key questions arise:

- What is the impact of newly resolved atmospheric variability on submesoscale eddies? Does this variability enhance or damp submesoscale activity?
- Which compensating mechanisms maintain observed variability levels?
- Should (and perhaps how) these interactions be parameterized in climate models?

To address these questions, the student will employ two approaches: (1) idealized submesoscale-resolving ocean simulations with kilometre-scale atmospheric forcing, and (2) state-of-the-art coupled model simulations to approximate real-world conditions.

Exploring the importance of multi-scale urban form, function and human behaviour dynamics to urban energy and water exchanges

Research question: Which processes need to be represented in kilometre scale Earth System simulations in order to capture the spatio-temporal variability of the urban microclimate?

The next generation of high-resolution urban modelling systems target a dynamic and multi-scale representation of feedback mechanisms between climate, form (building geometry, urban facets, materials), function (building use, infrastructure and systems), and human activities. This allows detailed process studies across spatio-temporal scales (people → building→ city; minutes→ seasons) that can provide the basis for enhanced parametrisations of urban signals in kilometre scale global models.

The project will use the land-surface model SUEWS-with Spartacus-Surface that represents the complexity of the urban energy balance from neighbourhood to building scale and vertically resolves radiative exchange processes and atmospheric state variables in the roughness sublayer. Coupled to SUEWS, the building energy model STEBBS provides building net heat storage, facet surface temperatures (roofs, walls) and indoor temperatures. Dynamics of anthropogenic heat emissions are captured by the agent-based model DAVE that represents human activities in indoor/outdoor microenvironments and associated energy consumption. This provides dynamic changes of building occupancy and internal heat gains to STEBBS (impacts anthropogenic building heat emissions). DAVE's modelling of the movement of citizens between urban neighbourhoods by various modes of transport (road, rail) allows to derive traffic-related heat emissions for SUEWS.

The project will apply the SUEWS-DAVE modelling system to different cities (e.g., London, Bristol, Berlin) to address the research question, including identifying which (and how) relevant urban descriptors can be derived from available global data.

Making digital rivers flow in km scale Earth System Modelling

Climate change, population growth and human intervention within river catchments make it more challenging than ever to provide reliable information on the current and future state of water in the world's rivers. Yet society urgently needs timely, dependable information for early warning of floods and droughts, which still today surprise and devastate communities.

Weather forecasting centres around the world use a global Earth System modelling approach to represent the interactions between the land, ocean and atmosphere to get the best forecasts of the upcoming weather and to forecast river flow. However, land surface hydrological processes are still not well represented in these models, particularly at km scale resolution and they suffer from substantial errors in predictions of infiltration, soil moisture, river flow, flood levels and energy/moisture feedbacks to the atmosphere. Yet these km scales are exactly those on which extreme rainfall, floods and the relevant preparedness actions take place.

The student will be embedded within the ECLand forecasting and research teams at ECMWF. This project will use global land-atmosphere coupled simulation experiments to assess how configuration of rainfall runoff and river discharge routing processes impacts the representation of surface fluxes. Experiments will test different land-atmosphere coupling configurations and parameterisations at the km scale, as well as ECMWF's machine learning hydrological predictions with the ultimate aim of improving predictions of river flow, soil moisture and feedbacks to operational weather forecasts.

Physical Models of Climate with Machine Learned Turbulent Fluxes

Machine learning in weather prediction is an ongoing revolution, creating accurate forecasts for a fraction of the cost of physical models (once trained). However, several limitations have thus far made pure data-driven models problematic for climate prediction:

1. Machine learned models do not conserve mass or energy, allowing the modelled climate's temperature to drift spuriously over time.

2. Machine learned models are not necessarily stable for long simulations – they can become increasingly unrealistic and eventually crash if they are run for longer than a few months.

3. Machine learned models cannot reliably make predictions outside their training data, limiting their use for projecting warmer climates.

We propose a different approach: rather than learning how atmospheric states evolve over time, we will develop a model that learns the fluxes -- how mass, energy and momentum are transported between regions. This approach guarantees conservation by design. By dividing the atmosphere into grid boxes and predicting changes based on what enters and leaves each box, we ensure that what leaves one box exactly enters another -- no mass or energy can be lost or gained through inaccurate modelling. Machine learning predicts the inter-box transports rather than the grid box values directly.

The aim is to produce an efficient and accurate model that 1) runs at lower resolution than purely physical models, 2) is cheaper to train than end-to-end machine learned weather models, 3) conserves fundamental properties like mass and energy by construction and 4) can predict new climate states and extreme weather.

Using XAI to inform future physical model development/improvement

The emergence of AI-based weather prediction systems, which can equal or surpass the predictive power of traditional numerical models, offers new opportunities for scientific advancement. Unlike conventional models that rely on explicitly formulated physical laws and parameterisations, AI models learn complex relationships directly from large observational and reanalysis datasets. This data-driven approach enables AI to capture processes, or handle complex error structures, that are incompletely represented in current physical frameworks.

The central research question of this PhD project is: "Can AI models be used to inform the improvement of physical models?" To address this, the student will explore how explainable artificial intelligence (XAI) methods can be applied to extract physically meaningful structures and interactions embedded within AI models. By systematically interpreting AI-driven predictions with XAI, the project aims to uncover insights into atmospheric processes that can guide the development and refinement of physical model components. Such insights could enhance model fidelity, improve predictive skill, and reduce biases, particularly in poorly constrained regimes such as extreme weather events or sub-grid-scale processes.

The research will investigate a range of XAI techniques, including feature attribution, layer-wise relevance analysis, and symbolic regression, to interrogate AI models and reveal the processes and relationships they capture. A key objective is to determine how these insights from XAI methods can translate into tangible improvements in physical models (for example, by suggesting refinements to parameterisations, identifying missing feedbacks, or revealing emergent process-level interactions).

By treating AI as both a predictive and exploratory tool, this project aims to deepen our understanding of which aspects of physical systems AI represents effectively. Ultimately, the research seeks to bridge the gap between data-driven and physics-based modelling, contributing to the next generation of weather prediction systems and enhancing our ability to forecast extreme or poorly understood atmospheric phenomena.

How does model resolution influence our understanding of the potential for iron-salt aerosol to deliver methane removal and marine cloud brightening?

Global Earth System models (ESMs) are central to the quantification of historical climate change and prediction of future climate change under different societal development scenarios. They are also being used ever more frequently to simulate potential climate change mitigation strategies such as stratospheric aerosol injection and marine cloud brightening, with the results forming an evidence base for policy makers.

However, while useful, ESMs suffer from biases when compared to present day observations, raising concerns about their ability to accurately simulate mitigation strategies both in terms of the potential cooling a strategy might deliver, but also a strategy's unintended consequences which could cause harm. A source of bias is the relatively low resolution (grid cells ~100 km wide) at which ESMs typically operate. This poses challenges for modelling processes which vary over small length scales, for example aerosol and chemistry processes stemming from point source emissions common with mitigation strategies.

We propose to run parallel simulations in the state-of-the-art United Kingdom Earth System Model (UKESM) at resolutions from the standard coarse approach down to the kilometre-scale. We will simulate injection of iron salt aerosol (ISA) into the marine atmosphere. ISA injection has been proposed as a way to remove atmospheric methane and simultaneously brighten marine clouds, potentially providing two routes for climate change mitigation. However, there remains significant uncertainty regarding ISA injection's impact and, given the small length scales over which realistic deployment would operate, simulating this strategy at high resolution is critical to assess its true potential.

Unpicking the Processes Leading to Hazardous European Windstorms

The British Isles are located at the end of the North Atlantic stormtrack and are frequently battered by cyclones, such as the recent storm Éowyn. The goal of this project is to understand and model the atmosphere and hazards at the scales where impacts of winds are felt, from hundreds of kilometres down to tens of metres. The new generation of km-scale coupled environmental prediction models will be confronted with new observations that have matching high-resolution spatial coverage.

The international NAWDIC field campaign, in early 2026, presents a major opportunity to deliver these observations. Aircraft will be deployed from Ireland with ground observations spanning Ireland, the UK and France. A key focus of NAWDIC is descending airstreams within cyclones and their interaction with the surface. The question is how momentum is transferred downwards to the surface, generating damaging wind gusts. This aspect is not well understood or modelled, with consequences for weather forecasting. Work using initial aircraft and buoy data indicates that extreme winds and ocean waves are underestimated. Changing surface drag relationships can improve one or the other, but not both, implying issues with atmospheric momentum transfer.

The student will exploit new Met Office and ECMWF model configurations to investigate the mechanisms linking windstorm processes with severe surface impacts, using the campaign observations to evaluate the most realistic configurations. There is scope to compare models with different architectures (e.g. testing LFRic, the next generation Met Office model) and to investigate whether AI models can predict the windstorm structure associated with the impacts.

Is representing fine scale phenomena in the Arctic crucial to estimate the resilience of AMOC?

The Atlantic Meridional Overturning Circulation (AMOC) is a key part of the climate system through its role in transporting heat. In the future, AMOC is expected to weaken due to CO2 emissions, and there is concern that AMOC could even decrease abruptly or irreversibly. However, the evidence for such projected AMOC decline originates from low resolution coupled climate models that have substantial deficiencies in their simulation of some important processes. One major area of uncertainty is the role of the Arctic in the formation of the dense 'head waters' of the AMOC. Many important high latitude processes such as boundary currents and mesoscale mixing are poorly resolved in low resolution models. Recent results indicate that the observed rate of dense water transformation in the Arctic may have increased in the last 30 years, making the AMOC more resilient than previously thought. However, whether similar changes are happening in models, and how these changes might be affected by the fundamental limitations of low-resolution models, is currently unknown.

To make progress, this project will evaluate the role of Arctic dense water transformation on the AMOC and understand how these will change in the future. The research will be based on the analysis of state-of-the art climate simulations using the Met Office coupled climate model. In particular, we will aim to understand whether the Arctic dense water formation is sensitive to the resolution of the model and characterize the key processes leading to any differences

What role do ocean eddies play in maintaining marine heat waves and how do they affect marine heat wave prediction?

Marine heatwaves are periods of extreme warm ocean temperatures that can have important impacts on weather and climate, as well as ecosystems and economies, in particular fishery closures resulting in food shortages. Therefore, it is important to understand what causes and maintains marine heatwaves and to be able to predict them in order to reduce their impacts. The project aims to better understand the role of mesoscale ocean eddies in maintaining marine heat waves and how the representation of mesoscale ocean eddies in ocean forecast models affects their skill in predicting marine heat waves on sub-seasonal to seasonal timescales.

The main research questions are:

1. What is the role of mesoscale ocean eddies in maintaining global surface and sub-surface marine heat waves?
2. What are the important eddy properties (e.g., number, eddy amplitude, transports) for marine heat waves on sub-seasonal and seasonal timescales?
3. What is the impact of not resolving ocean eddies in simulating and predicting marine heat waves in ocean only or coupled models?
4. How do eddy resolving models improve the representation and prediction of marine heat waves?

The project will use a broad range of global data sources from satellite altimeters, ocean reanalyses, high resolution ocean models and ocean forecasting systems (e.g. Mercator). To better understand the role of mesoscale ocean eddies in maintaining marine heat waves, the student will use a combination of Eulerian and

Lagrangian analysis techniques that include the use of an ocean heat wave index and direct tracking of ocean eddies in data sets that support them.

Improving subseasonal-to-seasonal forecasts in the tropics using AI

How can we use AIWPs to improve S2S ensemble forecasts in the tropics?

Subseasonal-to-seasonal forecasts (S2S; lead times of 15–60 days) are important in tropical countries for decision-making in agricultural, health, energy, and water security sectors. Current operational S2S forecasts are built using physics-based numerical weather prediction models (NWPs), and usually comprise either a single-model or multi-model ensemble that generates probabilistic forecasts. However, this approach leads to variable and often poor skill over many tropical regions.

AI-based weather prediction models (AIWPs) use a range of state-of-the-art machine learning models, trained on vast amounts of NWP-processed historical weather data, known as reanalyses, learning relationships without explicitly solving the underlying physical equations. As a result, AIWPs can produce weather forecasts with comparable skill to NWPs in large-scale fields but at a fraction of the cost. Some forecasting centres have thus started using them as a complementary addition to their existing NWP forecasts.

In this project, you will investigate the strengths and weaknesses of three types of AIWP at S2S lead times compared to traditional NWP based models: standard AIWPs (such as FourCastNet, trained to forecast up to ten days ahead), S2S AIWPs (such as FuXi-S2S, trained to forecast at S2S lead times), and hybrid models (such as NeuralGCM, which has an NWP at the core, but uses AI for parameterisations). You will then investigate how these are best integrated into existing ensemble prediction frameworks. This may take the form of a skill-weighted average, or an adaptive weighting based on regime-dependent skill.

Storm track regimes and medium range predictability.

The mid-latitude storm track is essentially a breeding ground for mid-latitude storms: strong horizontal temperature gradients on the western side of the ocean basins provide the required energy reservoir for midlatitude storms to grow. Each storm depletes some, or perhaps all of this energy reservoir. After such a depletion event, the energy reservoir needs building up again. Recent work has shown that the energy reservoir and the storm-track activity are related like a nonlinear predator-prey model.

The storm track and the jet stream form a complex interlinked system that is ultimately steered by this predator-prey system of storms feeding on upstream temperature gradients. It is expected that the dynamical core of weather forecast models should do well for such processes, yet some persistent model biases remain: the mean jet latitudes are often biased in a way that is consistent with too little upstream storm track activity. Furthermore, onset and decay of large-amplitude patterns, such as blocking patterns, are still notoriously hard to predict, thus degrading medium range forecast skill.

How good are current models in following the observed chain of processes from the predator-prey stage of nascent storms to the decay phase along the highly deformed jet stream? To what extent does this chain of processes rely on accurate representation of diabatic processes in the storm-track? Can we build this dynamical-system based framework for storm track variability into a physics-informed machine learning system to enhance medium range predictability?

Advancing Scalable and Interpretable Hybrid Modelling Frameworks for Kilometre-Scale Earth System Prediction

The transition to kilometre-scale Earth system models presents unprecedented opportunities for capturing fine-scale processes such as convection, urban dynamics, and land-atmosphere interactions. However, current state-of-the-art models - both physical and AI-driven including foundation models like Aurora and GraphCast - face significant limitations in scalability, interpretability, and generalisation across domains. Furthermore, extreme weather patterns (eg., peak temperatures in heatwaves, or intense wind speeds in storms like Ciaran) are difficult to capture. These could potentially benefit from physics-informed or causal AI-based emulators.

Hybrid modelling frameworks, integrating physical with data-driven components, offer a promising pathway to overcome these challenges. Yet, their development remains nascent, particularly at kilometre-scale resolution where computational demands, data sparsity, and model complexity converge.

Research Questions

- How can hybrid models be designed to scale efficiently across high-resolution domains without prohibitive computational cost?
- What techniques can ensure interpretability of AI components within hybrid systems?
- How transferable are hybrid models across different climate regimes and modelling systems?
- What are the trade-offs between accuracy, interpretability, and computational efficiency in hybrid kilometre-scale modelling?

This PhD will address the following key objectives:

1. Design scalable hybrid architectures operating efficiently at kilometre resolution, leveraging both physics-based and causal-AI components.
2. Develop interpretable modelling strategies enabling trust in hybrid predictions.
3. Establish transferability protocols to adapt hybrid models across different regions, climate regimes, and general circulation models (GCMs).
4. Benchmark hybrid frameworks against traditional models using real-world datasets and evaluation metrics relevant to Earth system prediction.

Developing the first sub-seasonal clear-air turbulence forecasts

Clear-air turbulence (CAT) is a significant hazard to aviation. Operational forecasts of CAT currently extend out to lead times of 18 hours and are used for flight planning. Projections show that climate change is strengthening CAT over periods of decades. However, there are currently no forecasts of CAT in between these two extremal time scales, despite demand from the aviation sector.

This project will design, implement, and test the world's first sub-seasonal forecasting system for CAT. The project will first develop a CAT forecast product for the IFS, by using a large multi-diagnostic ensemble to calculate eddy dissipation rates. A recent survey of ECMWF forecast users found that this product was high on their wish list. The project will then produce 12-hourly global probabilistic CAT forecasts out to four weeks using the 51 ensemble members of the SEAS5 forecasting system each month. The skill of the forecasts will be tested using automated in-flight eddy dissipation rate measurements from the WMO's Aircraft Meteorological Data Relay (AMDAR) system.

The end result will be a validated and verified CAT forecasting system suitable for operational use. It will allow probabilistic CAT forecasts out to four weeks to be produced and made available by ECMWF, updated monthly. It is known that there is demand for such a product from the aviation sector, which would use the forecasts for medium-term fuel strategy and for air-traffic controller and flight dispatcher workload planning. The product would also be suitable for inclusion in the next ECMWF climate reanalysis, ERA6.

A Machine Learning Framework for Multiscale Process Based Prediction in Earth System Models

Earth System Models (ESMs) are important tools that help us predict future climate events particular for impacts studies. However, ESMs are often coarse scaled due to computational constraints and do not capture subgrid scale processes essential for predicting regional impacts. Effective methods to downscale from coarser resolutions can allow us to run larger ensembles at relatively lower costs to obtain high resolution data.

In this project, we propose a machine learning (ML) framework to first statistically downscale key climate drivers at lower resolution (global 10km and global 5km) and then to evaluate the downscaled projections so we can better understand which processes require higher resolutions (1 km or finer) for improved regional impacts assessments. We propose to develop a temporally coupled Convolutional Neural Network based architecture for multivariate downscaling to answer the following key research questions:

1) How fit for purpose is the downscaled data prediction and how does its skill compare against higher resolution models? We will use observations to evaluate how well the system captures characteristic features of statistically sparse and small-scale convection systems such as rainfall intensity, duration and spread.
2) Can we gain an understanding of what subgrid scale process effects are effectively captured by the ML architecture? We propose to use an interpretable ML based approach comprising of feature importance ranking methods and emulators to investigate causality relationships. Such an approach will build trust in the downscaling and can help identify the underpinning physical processes.