

Thursday, April 23, 2026, 07-09 UTC



## Topic: Machine Learning in DA

**Organizers:** James Taylor (RIKEN, Japan),  
Michael Goodliff (RIKEN, Japan)

---

### Program:

**07:00 – 07:05**      **Welcome**

---

**07:05 – 07:25**      **AI-DOP: A machine-learned path from Earth System observations to a global weather forecast**  
(18' + 2')

*Mihai Alexe, Eulalie Boucher, Peter Lean, Ewan Pinnington, Simon Lang, Patrick Laloyaux, Tomas Kral, Tobias Necker, Tony McNally*

---

**07:25 – 07:45**      **Efficient ensemble forecasting with generative flow models and data assimilation**  
(18' + 2')

*Tobias Finn, Marc Bocquet, Freddy Bouchet, Etienne Mémin*

---

**07:45 – 08:05**      **Forward and inverse modelling of parametric dynamical system in a physics-aware latent space**  
(18' + 2')

*Qiyao Zhou, Xujia Zhu, Pierre Joli, Yu Cong, Sibor Cheng*

---

**08:05 – 08:25**      **Unifying background-error covariance modelling for midlatitude and tropical atmospheric data assimilation**  
(18' + 2')

*Bostjan Melinc, Uros Perkan, Lukas Kugler, Ziga Zaplotnik*

---

**08:25 – 08:45**      **High-Resolution Ensemble Uncertainty Quantification Using Diffusion-Based Machine Learning**  
(18' + 2')

*Swapan Mallick, Kasper Tølløse, Per Dahlgren, Jelena Bojarova, Xiaohua Yang, Eivind Støylen, Harald Schyberg*

---

**08:45 – 08:50**      **Closing**

The times in UTC are approximate.

In case of technical problems, we might have to change the order of the presentations.

### Time Zones: 07 – 09 UTC

Europe:	8 - 10 am BST (London)	9 – 11 am CEST (Berlin)	
Asia/Australia:	3 – 5 pm CST (Shanghai)	4 – 6 pm JST (Tokyo)	5 – 7 pm AEDT (Sydney)
Americas:	0 – 2 am PDT (San Fran.)	1 – 3 am MDT (Denver)	3 – 5 am EST (New York)

## **AI-DOP: A machine-learned path from Earth System observations to a global weather forecast**

**Mihai Alexe**<sup>1</sup>, Eulalie Boucher<sup>1</sup>, Peter Lean<sup>1</sup>, Ewan Pinnington<sup>1</sup>, Simon Lang<sup>1</sup>, Patrick Laloyaux<sup>1</sup>, Tomas Kral<sup>1</sup>, Tobias Necker<sup>1</sup>, Tony McNally<sup>1</sup>

<sup>1</sup>ECMWF

Recent advances in machine learning (ML) for numerical weather prediction (NWP) have produced models that match, or in some cases surpass, the forecast skill of leading physical models such as ECMWF's Integrated Forecast System (IFS) in several key metrics. Almost invariably, these models are trained on large NWP (re)analysis datasets that provide complete three-dimensional descriptions of the atmospheric state spanning multiple decades. A (re)analysis is traditionally produced using data assimilation (DA), a procedure that combines observations with a physical model trajectory to produce statistical best estimates of the atmosphere. An undeniable success story of NWP, state-of-the-art DA systems have become highly sophisticated, reflecting decades of scientific progress.

AI-Direct Observation Prediction (DOP) is an R&D project at ECMWF that aims to determine if it is possible to learn a skillful medium-range weather forecast directly from Earth System observations (e.g., satellite radiances and in-situ data), i.e., without input from traditional NWP. Current research results indicate that the DOP system learns a coherent internal representation of the Earth System and can generalize well to regions that are only sparsely observed. Moreover, DOP is able to capture cross-domain interactions – such as atmosphere-ocean - by learning to use information present in “interface” observations.

The talk will give an overview of the latest results from AI-DOP. We will also present recent progress made towards the integration of Earth System observation data into the award-winning open-source Anemoi ecosystem that is developed jointly at ECMWF and at many meteorological agencies across Europe. Finally, we will discuss potential pathways toward a real-time (operational-like) end-to-end data-driven forecasting system.

# Efficient ensemble forecasting with generative flow models and data assimilation

Tobias Finn<sup>1</sup>, Marc Bocquet<sup>1</sup>, Freddy Bouchet<sup>2</sup>, Etienne Memin<sup>3</sup>

<sup>1</sup>CEREA, Ecole des Ponts, IP Paris, IPSL

<sup>2</sup>LMD, Ecole normale supérieure

<sup>1</sup>Inria, Université de Rennes

Deep learning forecasting models have transformed weather prediction and Earth system modelling, often outperforming traditional physics-based models. However, when, these models are trained to produce single predictions, they face two key problems: they lose fine-scale details (resulting in overly smooth predictions) and they don't always respect physical laws. Generative flow models, which encapsulate diffusion models and flow matching, offer a solution by learning to progressively refining predictions across multiple scales. Using sea-ice modelling as an example, I will show how these models can efficiently generalise from short-term forecasts to decade-long simulations while maintaining physical consistency. Building on this foundation, I will explain how data assimilation can be integrated into generative forecasting models. Drawing on previous research, I will present a novel ensemble-based approach that efficiently leverages the generative model's structure to align predictions to observations.

# Forward and inverse modelling of parametric dynamical system in a physics-aware latent space

Qiyao Zhou<sup>1</sup>, Xujia Zhu<sup>2</sup>, Pierre Joli<sup>1</sup>, Yu Cong<sup>1</sup>, Sib0 Cheng<sup>3</sup>

<sup>1</sup>LMEE, UnivEvry, Universite Paris-Saclay, Evry, France

<sup>2</sup>L2S, Universite Paris-Saclay, CNRS, CentraleSupélec, Gif-sur-Yvette, France

<sup>3</sup>CEREA, ENPC, EDF R&D, Institut Polytechnique de Paris, Ile-de-France, France

Forward and inverse modeling in high-dimensional dynamical systems remain central challenges in computational fluid dynamics and data assimilation (DA). Although reduced-order models (ROMs) have significantly improved computational efficiency, their integration with DA often suffers from limited physical interpretability, weak parameter identifiability and reduced efficiency when solving inverse problem in real time. In particular, latent representations learned purely from data may not preserve meaningful correlations with underlying physical parameters, thereby limiting their effectiveness in inverse modeling tasks.

In this work, we propose a physics-aware neural-network-based variational DA scheme that systematically integrates deep-learning-based ROMs with variational DA for parameter calibration in high-dimensional dynamical systems. The main idea is to perform variational DA in a structured latent space that is learned to be explicitly correlated with underlying physical parameters through observable supervision. Conducting DA in this physics-aware latent space enhances regularity, improves parameter identifiability, and reduces uncertainty in both calibration and subsequent state prediction.

The proposed framework provides a unified and scalable strategy for coupling deep-learning-based surrogate modeling with variational DA in inverse problems. Numerical experiments conducted on two canonical fluid dynamics systems demonstrate that the physics-aware deep-learning-based framework significantly outperforms standard ROM-DA approaches in terms of robustness and parameter estimation accuracy, while preserving reliable forward predictive capability.

# Unifying background-error covariance modelling for midlatitude and tropical atmospheric data assimilation

**Bostjan Melinc**<sup>1</sup>, Uros Perkan<sup>1</sup>, Lukas Kugler<sup>1</sup>  
Ziga Zaplotnik<sup>1,2</sup>

<sup>1</sup> University of Ljubljana, Faculty of Mathematics and Physics

<sup>2</sup> ECMWF

Estimating and modelling background-error covariances poses a core challenge in variational data assimilation (DA). Operational systems typically approximate these covariances by a set of analytical and empirical transformations that separate geostrophically balanced components from unbalanced inertia-gravity modes - an approach well-suited for the midlatitudes but less applicable in the tropics, where different physical balances prevail. Here, we estimate background-error covariances in a reduced-dimension latent space learned by a convolutional autoencoder (AE). The AE was trained using 40 years of ERA5 reanalysis data, enabling it to capture flow-dependent atmospheric balances from a diverse set of weather states. We demonstrate that performing DA in the latent space yields analysis increments that preserve multivariate horizontal and vertical physical balances in both tropical and midlatitude atmosphere. Assimilating a single observation of 500 hPa geopotential height in the midlatitudes produces increments consistent with geostrophic and thermal wind balance, while assimilating a total column water vapour observation with a positive departure in the tropical atmosphere with nearly saturated humidity generates an increment resembling the tropical response to heat-induced perturbations. Furthermore, the resulting increments are localised, flow-dependent, and shaped by orography and land-sea contrasts. Finally, we explore the transition from using synthetic ensembles and a climatology-based background-error covariance matrix to an operational ensemble of data assimilations (EDA). Despite significant variance loss in some variables due to compression, latent-space assimilation retains balanced, flow-dependent increments. This type of background-error covariance modelling may therefore also be suitable for operational applications – both in a hybrid or a full machine-learning-based variational DA setup.

## Exploration of tempering data assimilation methods for using the LETKF: Simple and intermediate experiments

Swapan Mallick<sup>1</sup>, Kasper Tølløse<sup>2</sup>, Per Dahlgren<sup>3</sup>, Jelena Bojarova<sup>1</sup>, Xiaohua Yang<sup>2</sup>, Eivind Støylen<sup>3</sup>, Harald Schyberg<sup>3</sup>

<sup>1</sup> Swedish Meteorological and Hydrological Institute (SMHI)

<sup>2</sup> DMI

<sup>3</sup> Met-Norway

The Copernicus Arctic Regional Reanalysis version 2 (CARRA2), distinguished by its high spatial resolution, provides improved representation of atmospheric processes, particularly those influenced by orography, relative to coarser reanalysis products. CARRA2 assimilates an extensive observational dataset within the HARMONIE numerical weather prediction system, coupled with a three-dimensional variational data assimilation scheme, to generate a high-fidelity reconstruction of the atmospheric state. A key objective of CARRA2 is the quantification of ensemble-based uncertainty for essential climate variables. To address this objective, we develop a machine learning framework based on denoising diffusion probabilistic models to infer high-resolution uncertainty estimates by leveraging paired ensembles from ERA5 (62 km) and CARRA2 (2.5 km). Surface variables are trained in a supervised manner, using ERA5 ensemble fields as inputs and corresponding CARRA2 fields as targets. The diffusion model employs a U-Net-based super-resolution architecture, enabling dynamic, data-driven resampling within the diffusion process while preserving fine-scale spatial structures that are often degraded by traditional downscaling approaches. The resulting diffusion-based framework effectively captures the cross-scale uncertainty relationship between the coarse-resolution ERA5 ensemble and the high-resolution CARRA2 ensembles. It produces high-resolution gridded uncertainty fields that accurately represent flow-dependent and orographically induced variability with strong spatial coherence. These findings highlight the robustness and potential of diffusion-based methodologies for uncertainty quantification in regional climate reanalysis.