AI-DOP: Data-driven “Direct Observation Prediction”

… or, how to learn a skillful medium-range weather forecast only from Earth System observations

Mihai Alexe (on behalf of the AI-DOP group at ECMWF, with special thanks to Eulalie Boucher, Peter Lean and Tony McNally)

ECMWF Bonn

firstname(dot)lastname@ecmwf.int

The rationale

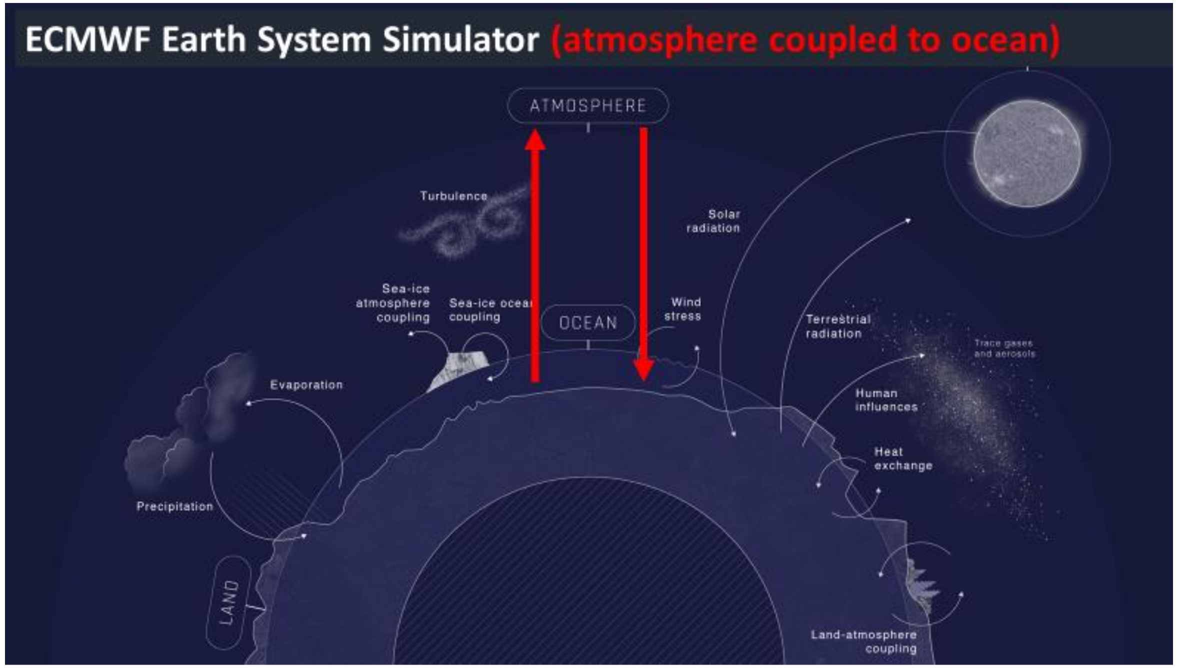

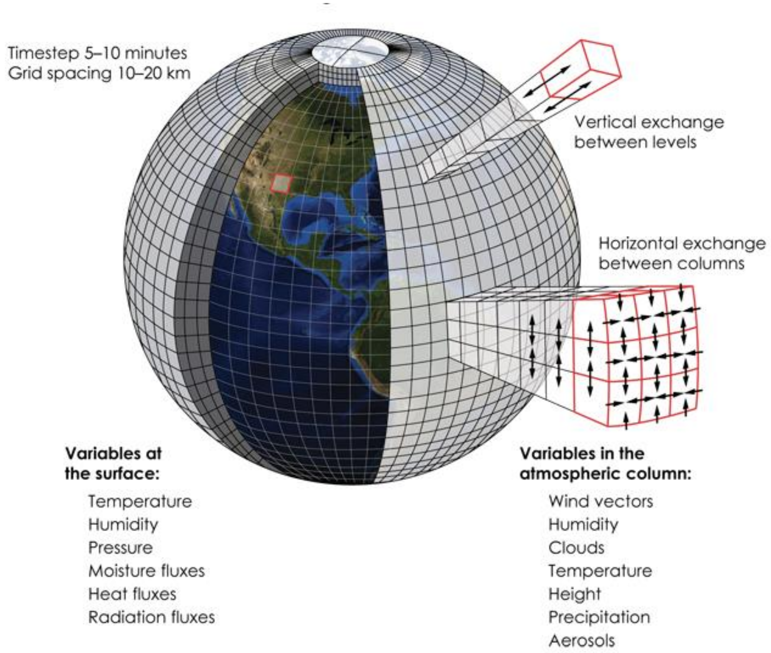

$$J(\delta\mathbf{x}_0) = \frac{1}{2}\delta\mathbf{x}_0^T\mathbf{B}^{-1}\delta\mathbf{x}_0 + \frac{1}{2}\sum_{i=0}^{n}(\mathbf{H}_i(\delta\mathbf{x}_i) - \mathbf{d}_i)^T\mathbf{R}_i^{-1}(\mathbf{H}_i(\delta\mathbf{x}_i) - \mathbf{d}_i)$$

$$\nabla_{\delta\mathbf{x}_0}J = \mathbf{B}^{-1}\delta\mathbf{x}_0 + \frac{1}{2}\sum_{i=0}^{n}\mathbf{M}^T(t_i, t_0)\,\mathbf{H}_i^T\,\mathbf{R}_i^{-1}\,(\mathbf{H}_i(\delta\mathbf{x}_i) - \mathbf{d}_i) = 0$$

Can we use the tools and models developed as part of the NWP ML revolution to bypass some of the scientific and technical challenges in state-of-the-art data assimilation systems?

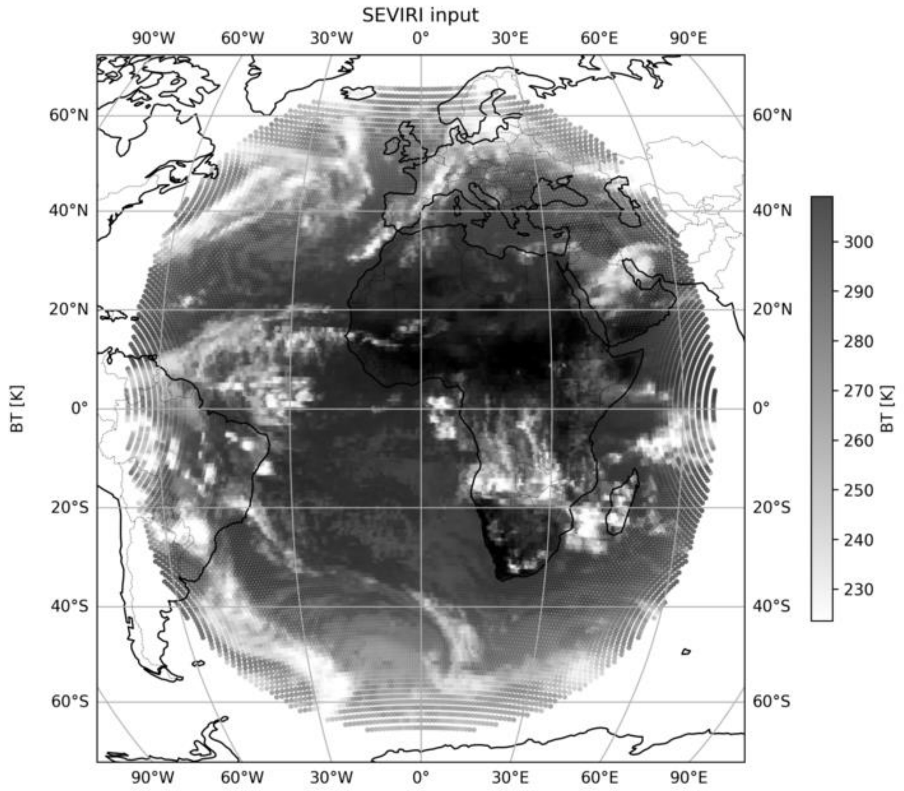

Conventional and satellite observations provide a rich amount of information about the atmospheric state and its spatio-temporal evolution …

so …

Can we learn a medium-range forecast directly from observations?

AIFS (Lang et al., 2024)

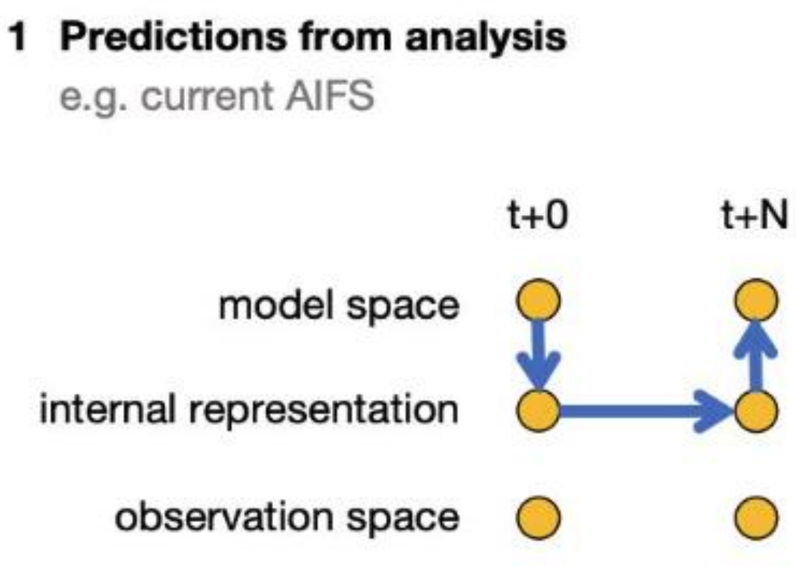

AI-DOP

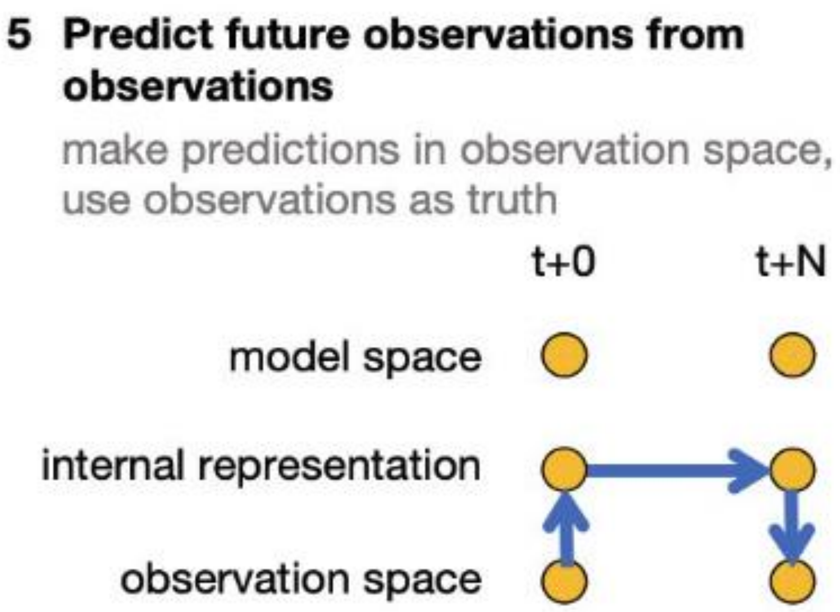

- **Use historical observations to train a neural network to forecast <u>future</u> observations (<u>don't need analyses</u>)**
- Include all available observations of the <u>full Earth system</u> (atmosphere, ocean, land) simultaneously
- Once trained, we initialize the model <u>directly</u> with the observations
- The model can produce a forecast at unobserved locations (e.g., on a grid)

ECMWF

The observations

Observations

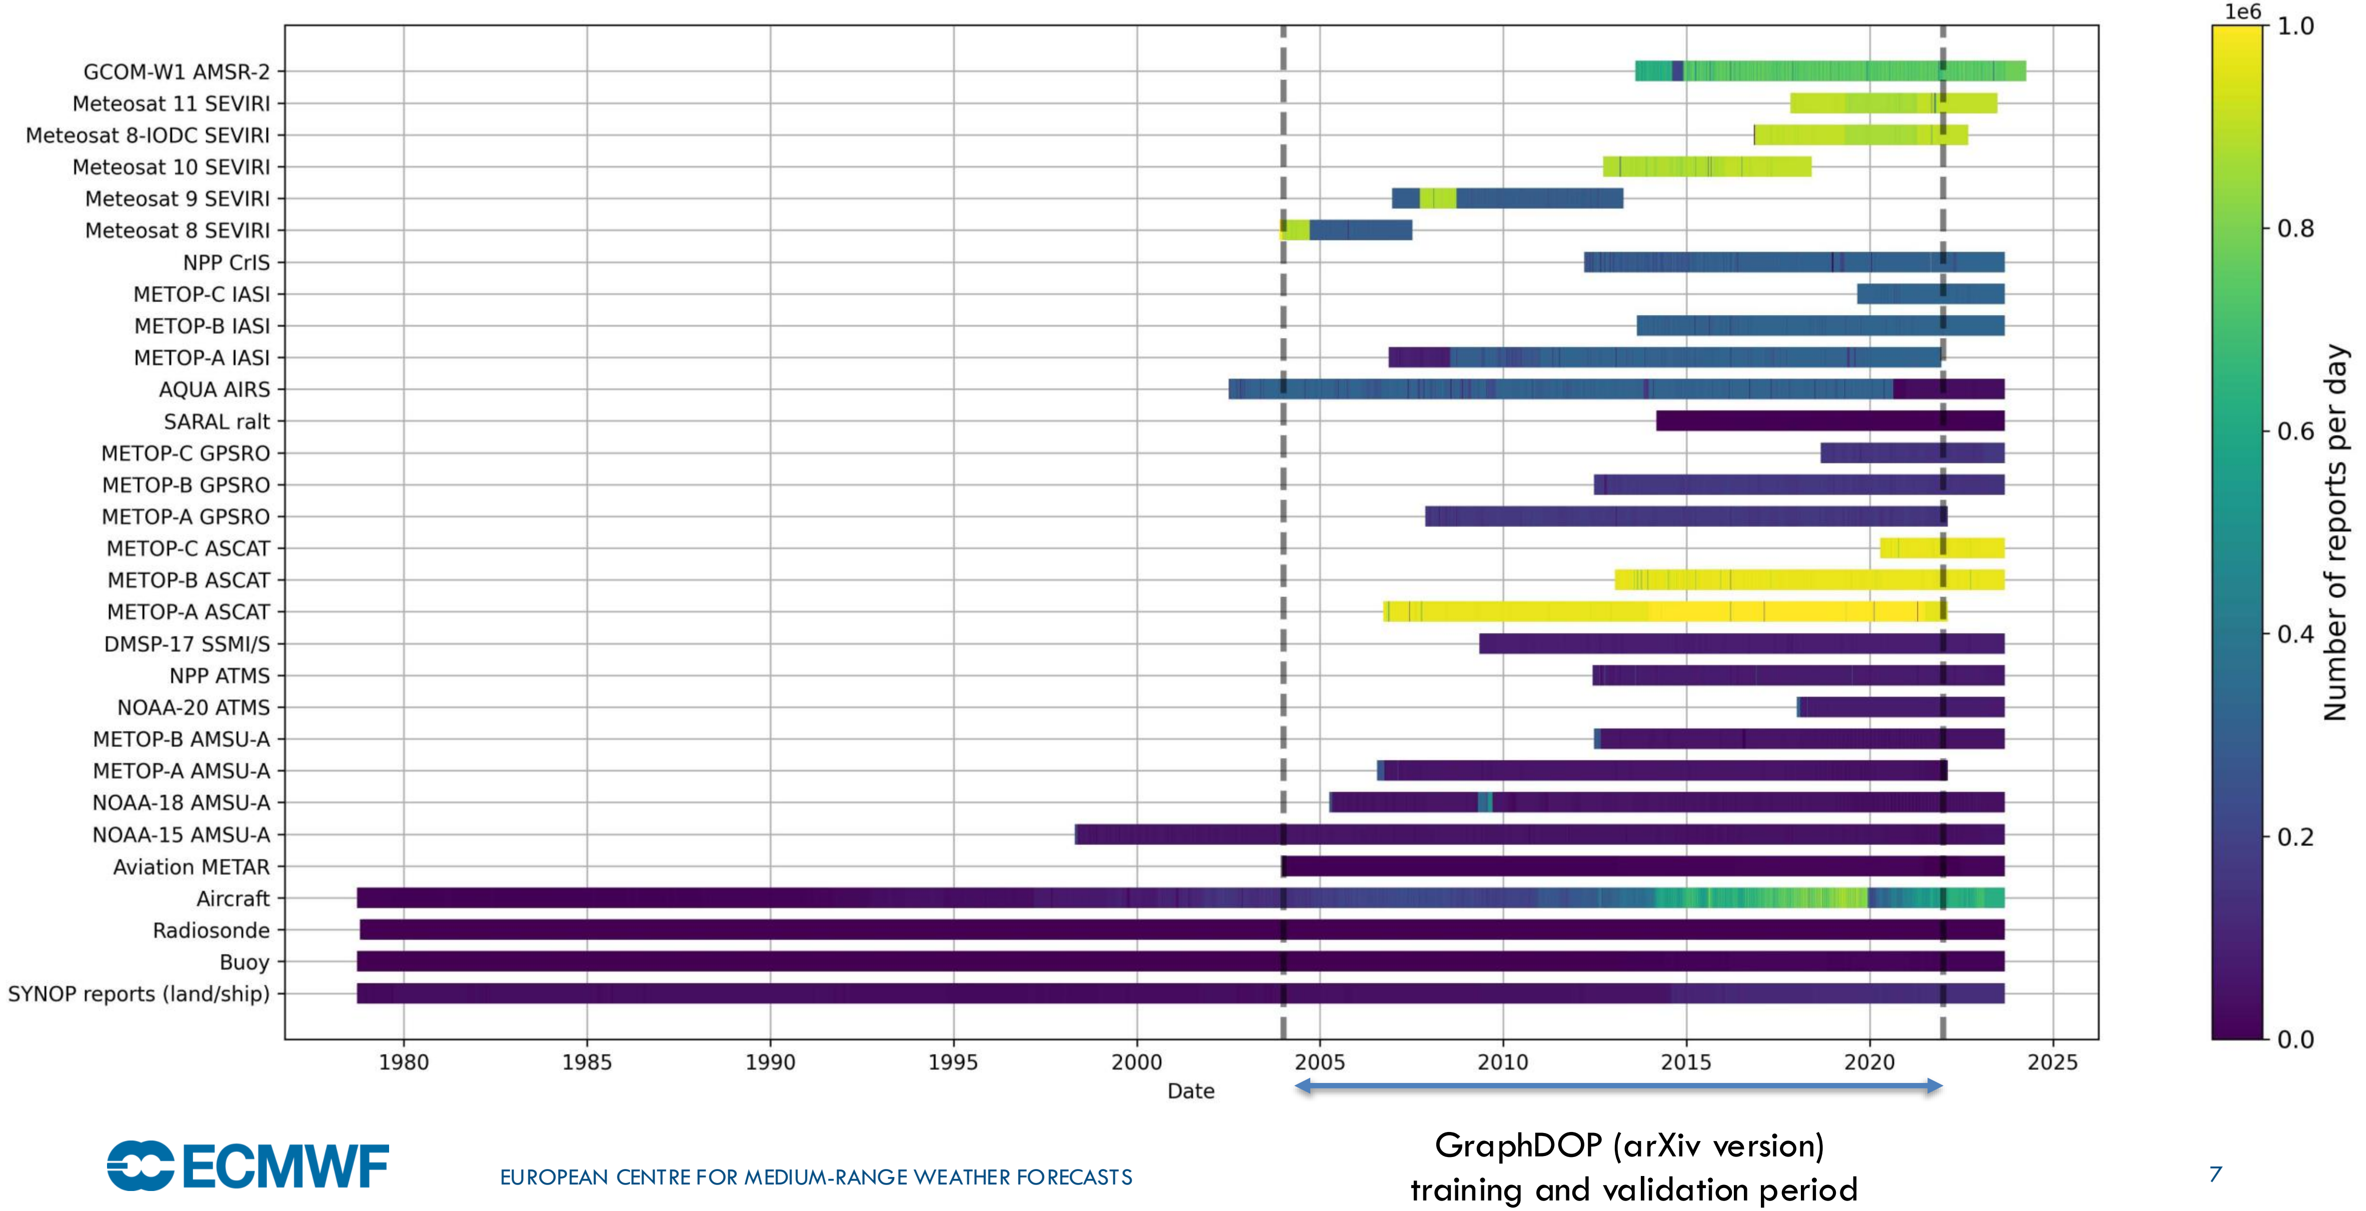

As of April 2025, ca. **81 billion reports** available for us to train on – ca. **7TB of zarr data**
More instruments are being added …

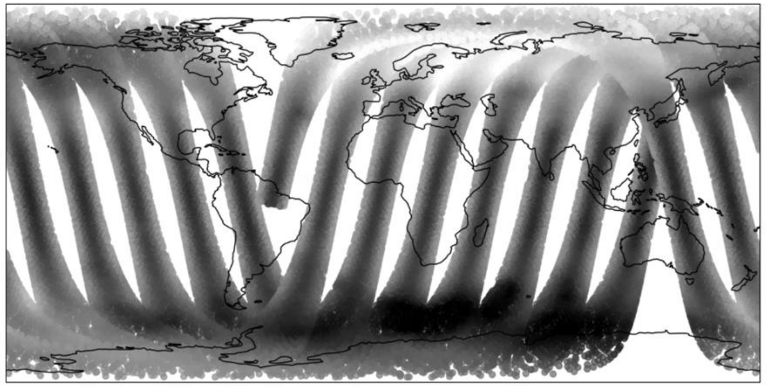

METOP-B AMSU-A ch7

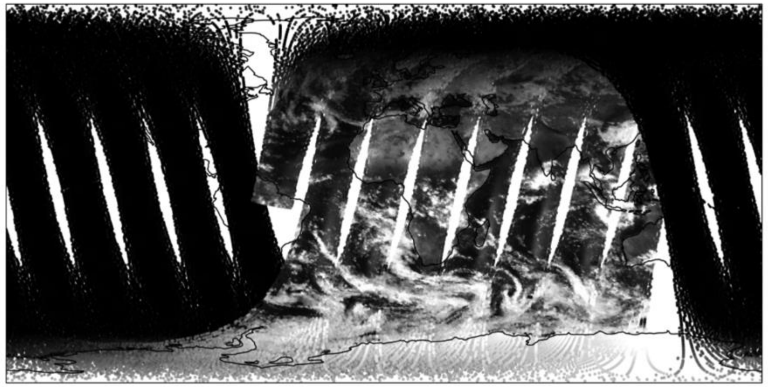

AVHRR visible reflectances

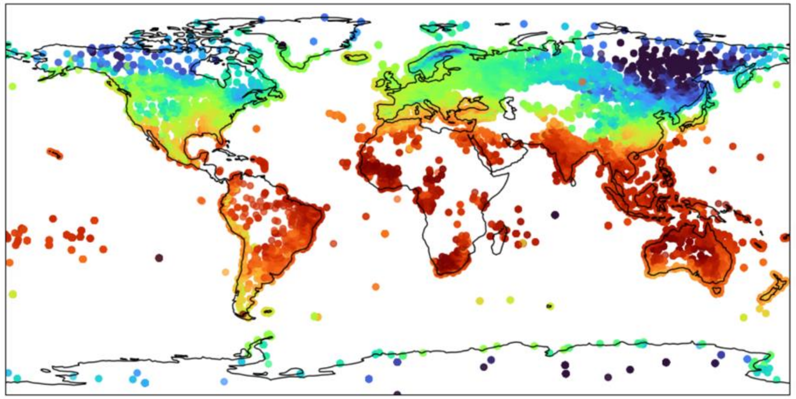

BUFR Land SYNOP 2m Temperature

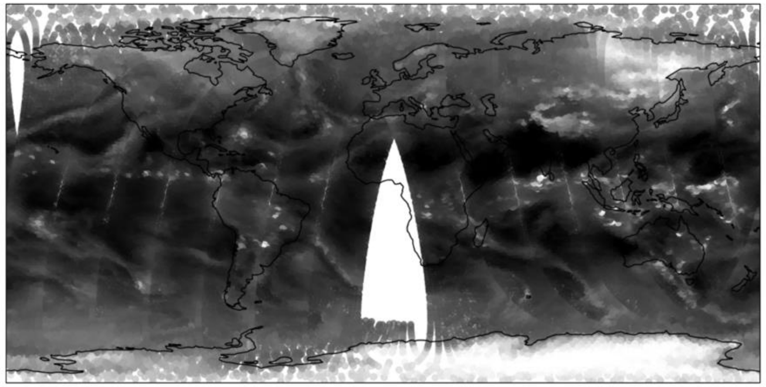

NPP ATMS ch18

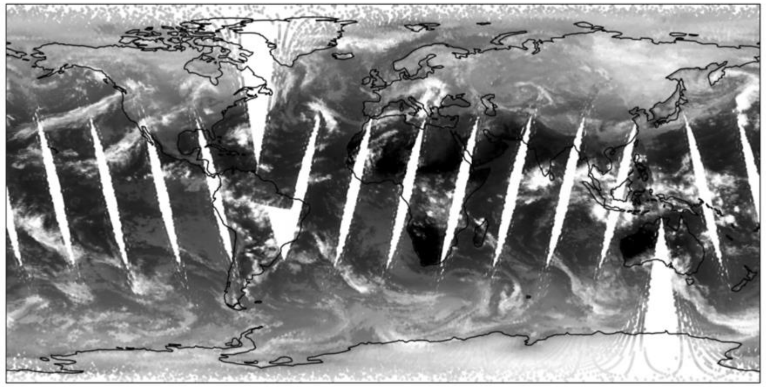

METOP-B IASI ch756

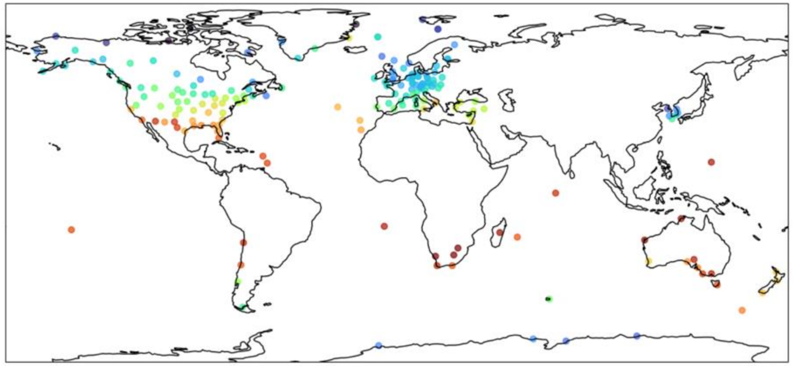

BUFR Land TEMP 850hPa Temperature

The neural network

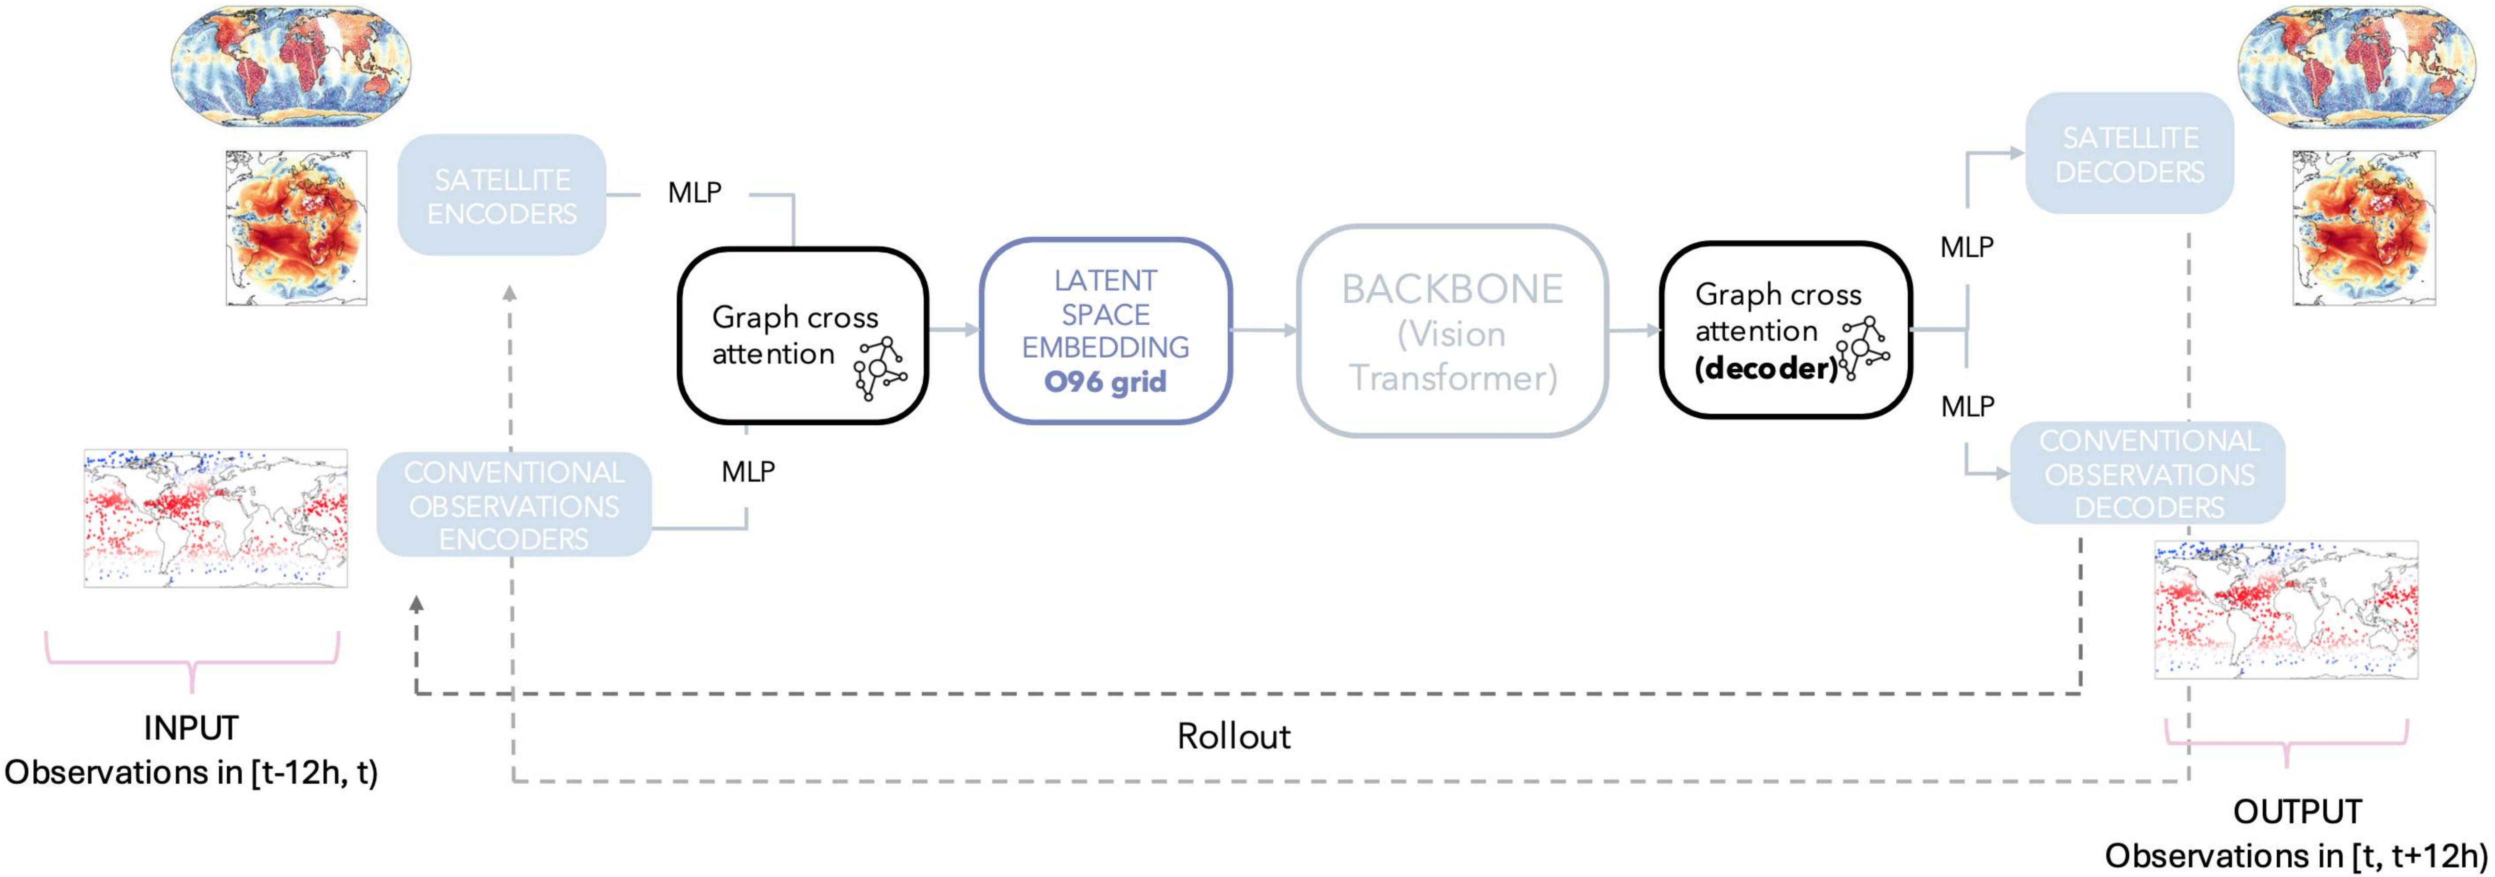
SATELLITE ENCODERS
MLP
Graph cross attention
LATENT SPACE EMBEDDING
O96 grid
BACKBONE (Vision Transformer)
Graph cross attention (decoder)
MLP
SATELLITE DECODERS
MLP
CONVENTIONAL OBSERVATIONS DECODERS
CONVENTIONAL OBSERVATIONS ENCODERS
MLP
INPUT
Observations in [t-12h, t)
Rollout
OUTPUT
Observations in [t, t+12h)

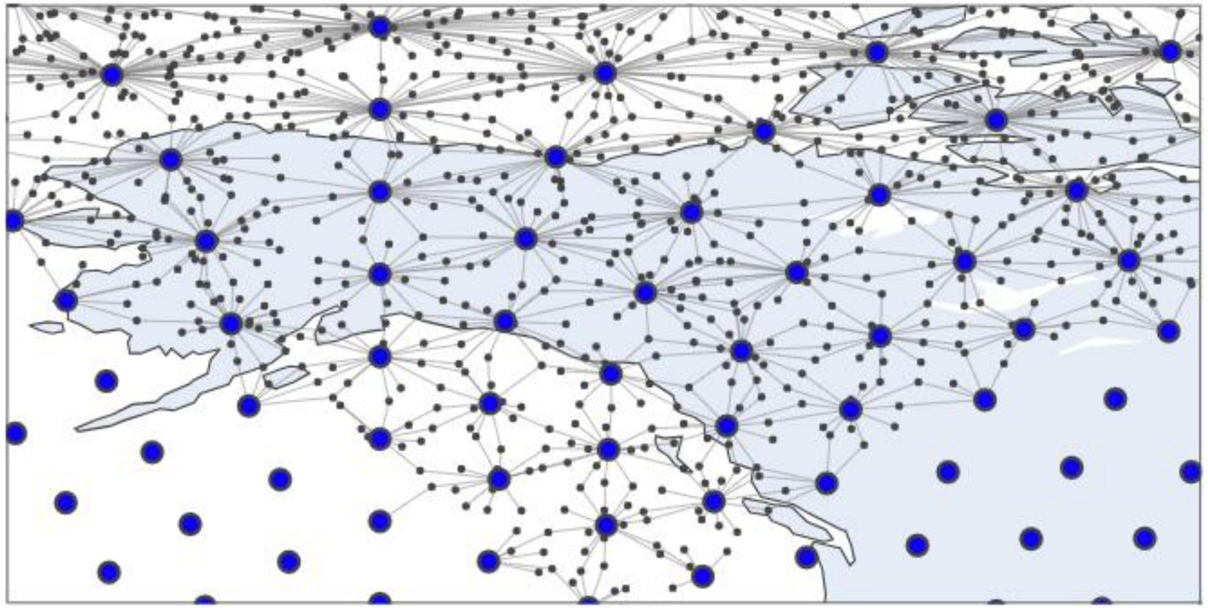

OBS-to-H graph

METOP-B AMSUA
channel 13 radiances

12-hour window

H to OBS graph

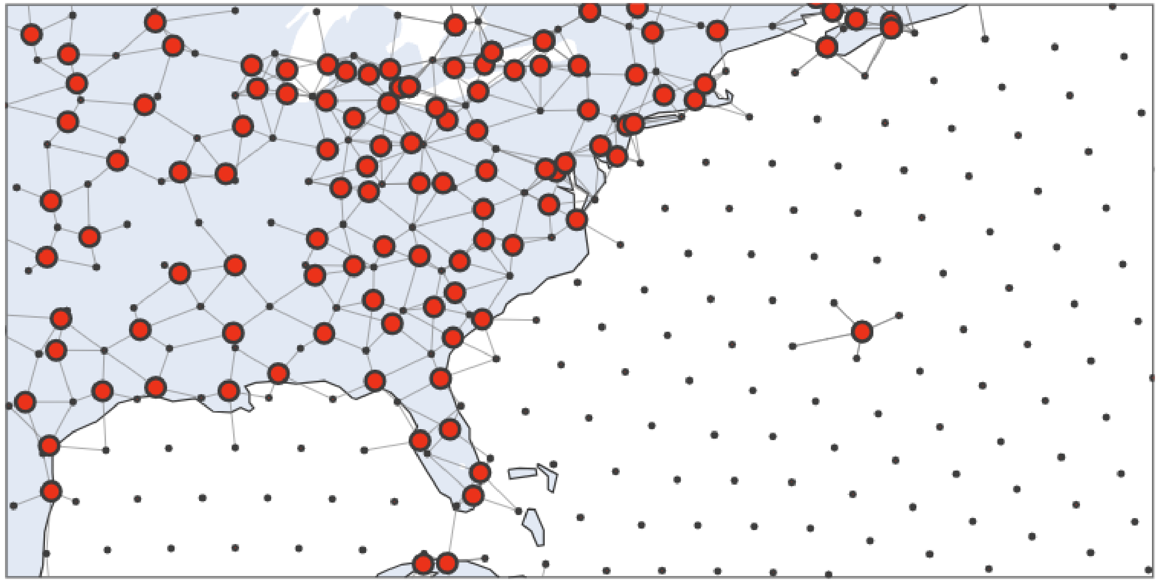

SYNOP stations
(12-hour window)

ECMWF

Does this work?

https://arxiv.org/abs/2412.15687

SEVIRI prediction | SEVIRI target

10.8μm

6.2μm

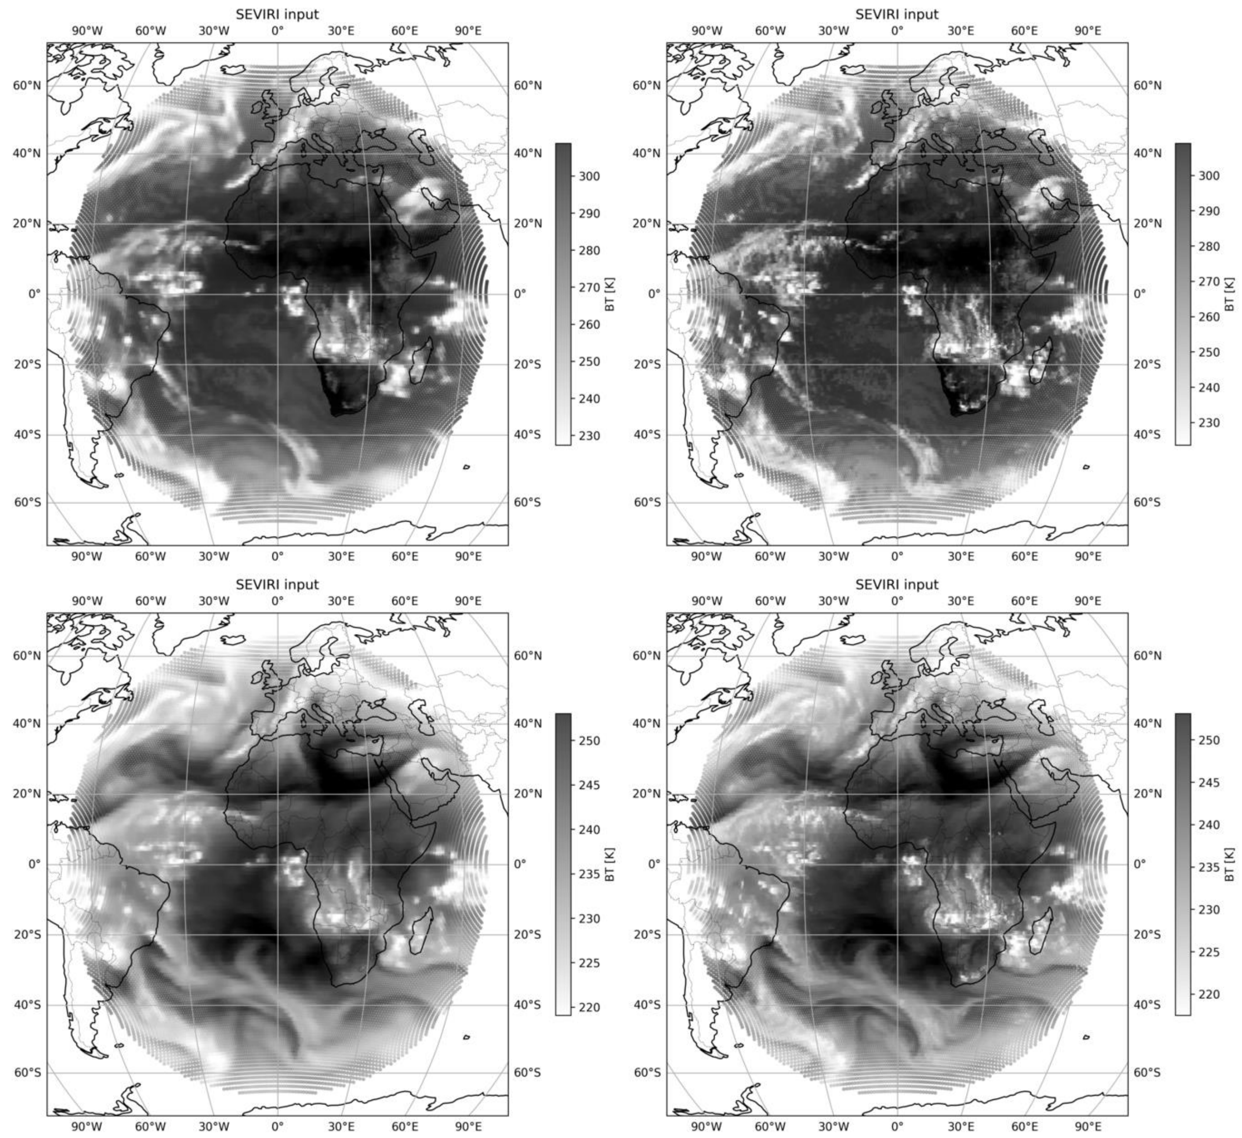

ECMWF

Forecasts of satellite observations

SEVIRI water-vapour 6.2 channel 5

AI-DOP | Target

day 1

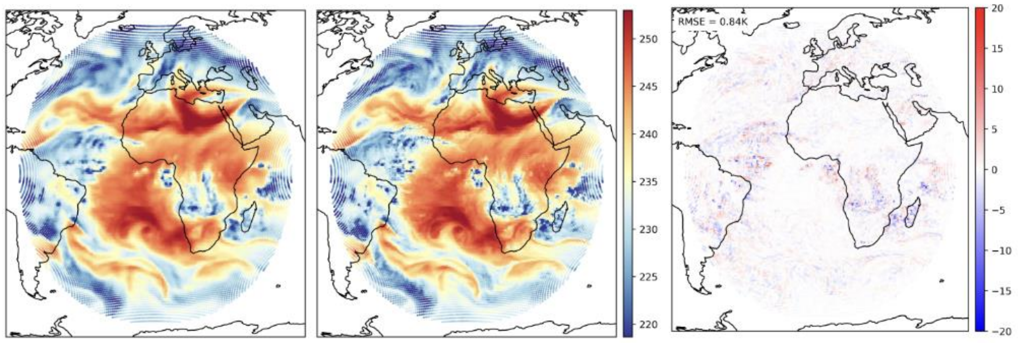

(a) t+10h (Jan 2, 2023, 09:45z) 45 minutes after the end of the input window

day 2

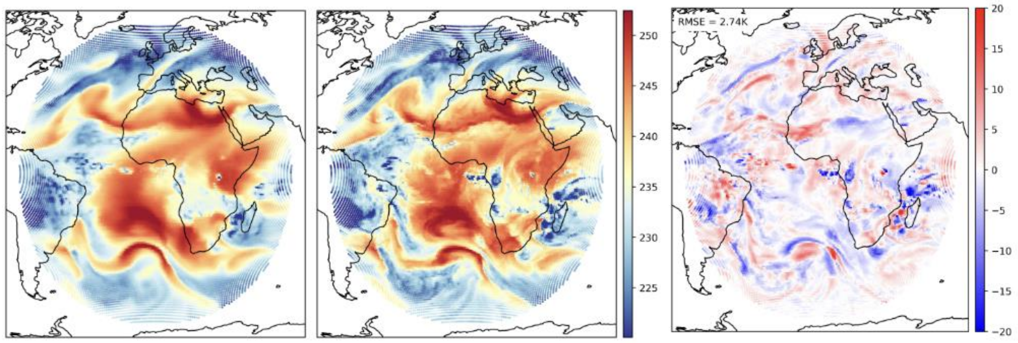

(b) t+33h (Jan 3, 2023, 08:45z)

day 3

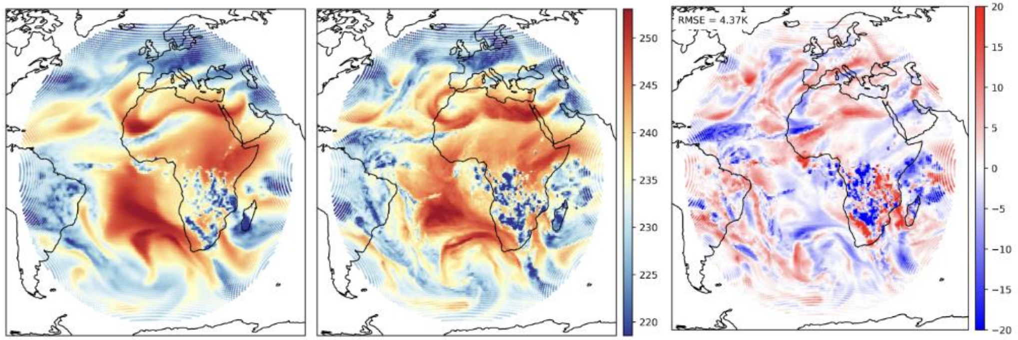

(c) t+63h (Jan 4, 2023, 14:45z)

IASI (window) channel 921 (875cm-1)

AI-DOP | Target

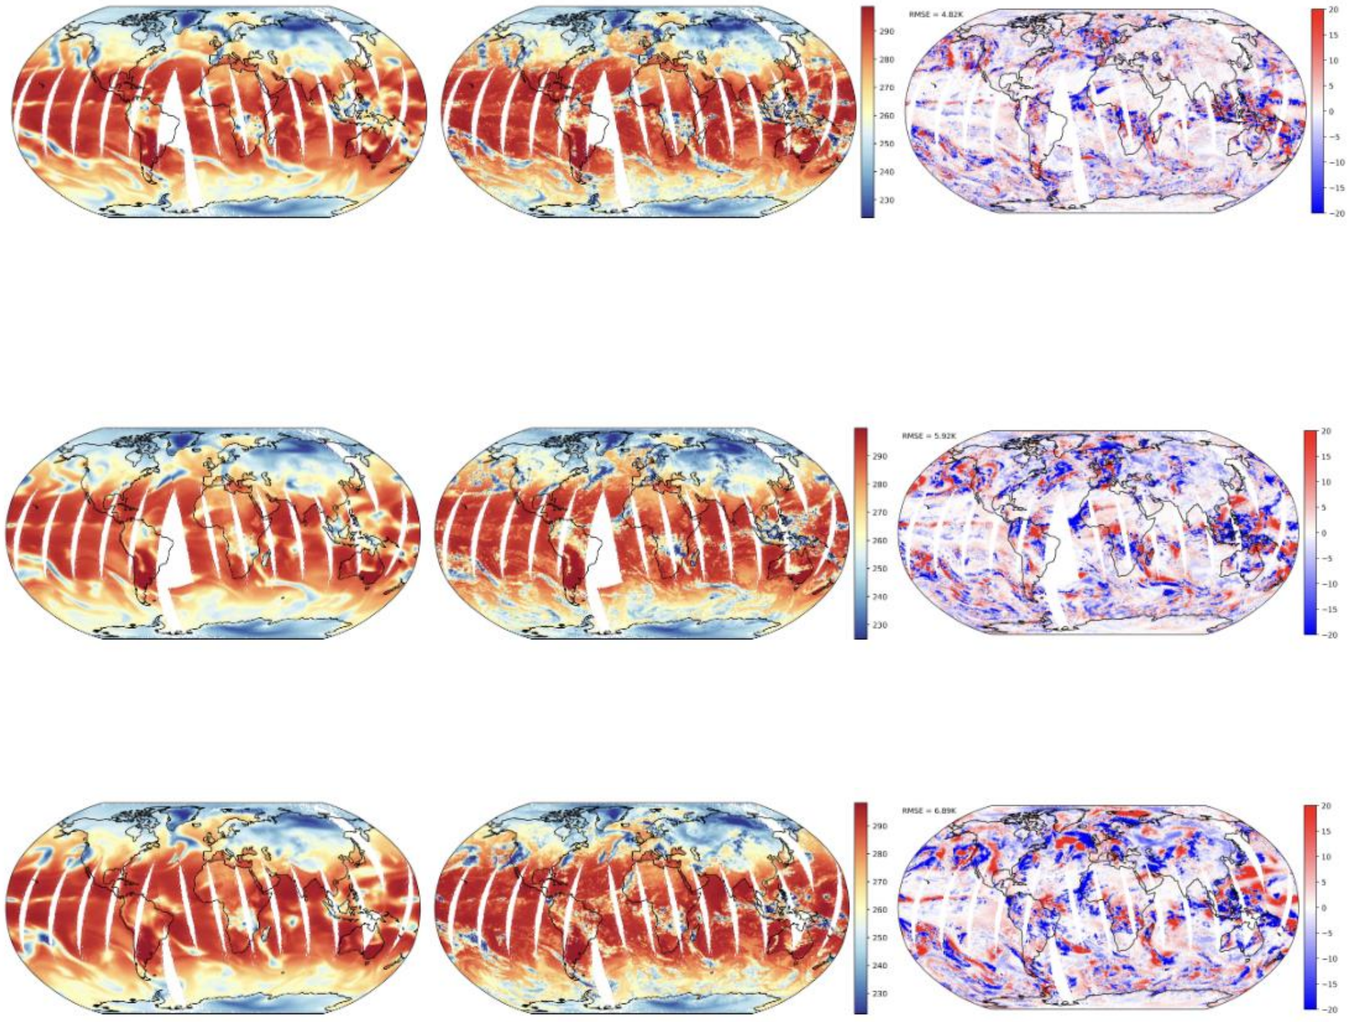

GraphDOP can forecast observations on a **regular grid** (here o96)

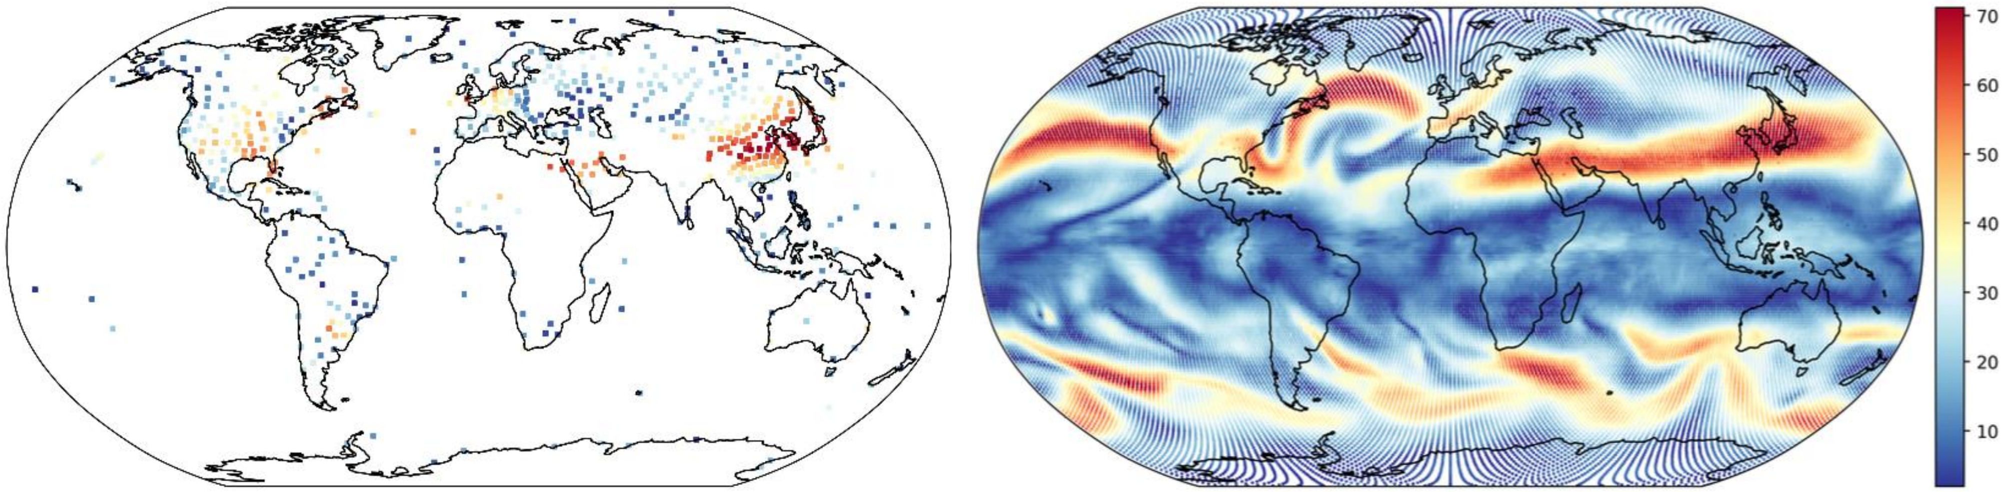

Figure 6: Observations of wind speed at 200 hPa used as input to the network (left) and the gridded 200 hPa wind speed from a 24-hour GraphDOP forecast (right) valid on Jan 15, 2023, 12z.

Gridded forecasts

AI-DOP ERA5

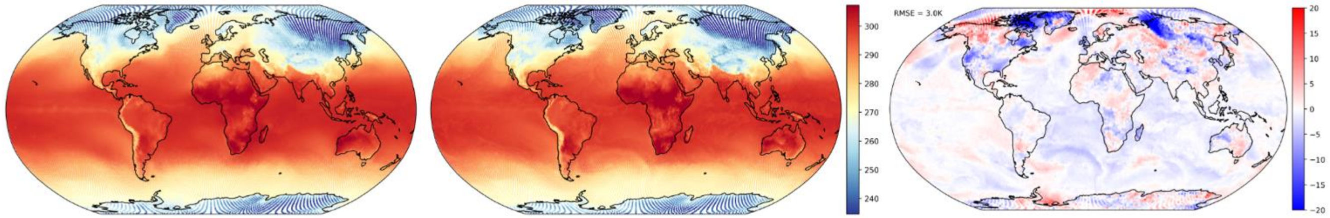

2m temperature

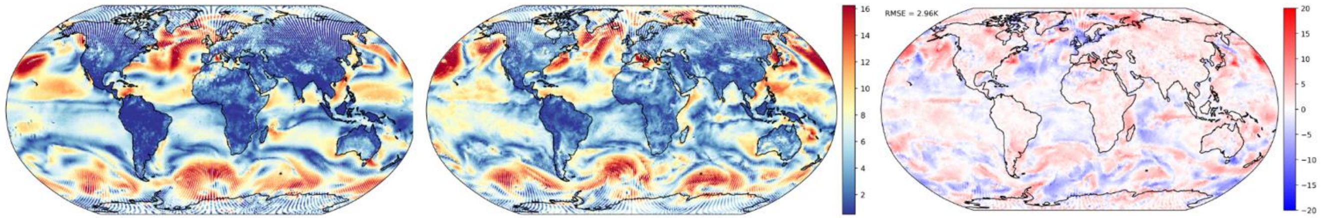

10m wind speed

AI-DOP ERA5

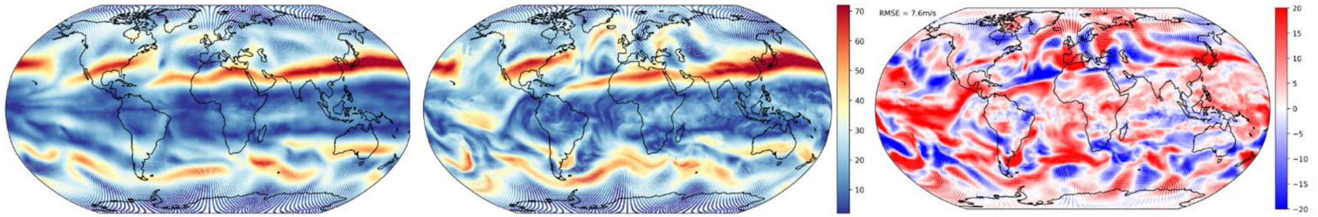

Wind speed 200hPa

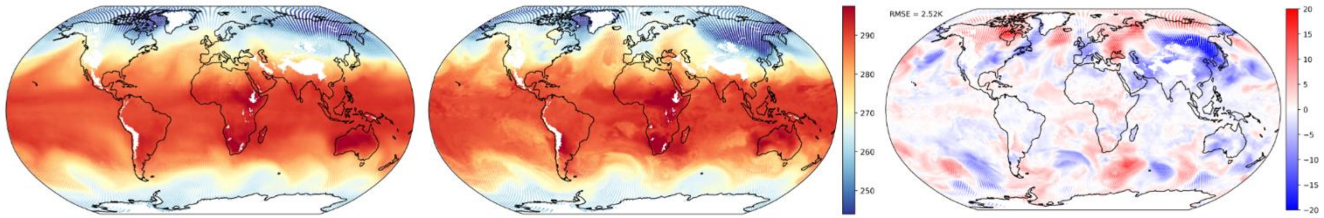

Temperature 850hPa

ECMWF

How well does it work?

Weather parameters (t2m) at observation locations

- Globally, slightly better than IFS at 24h, slightly worse at 120h
- Better than IFS over the tropics out to 120h (and beyond)
- Winter high latitude (snow) surfaces clearly a problem

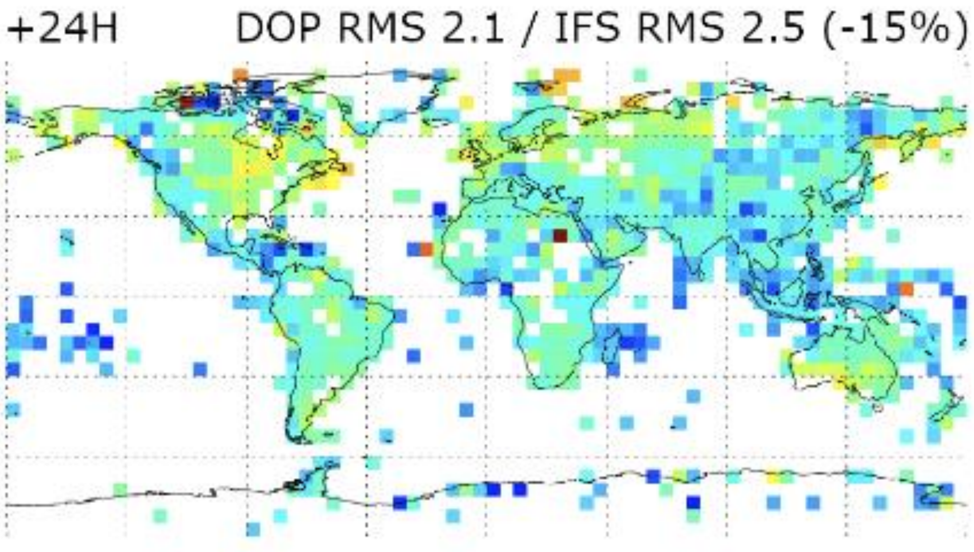

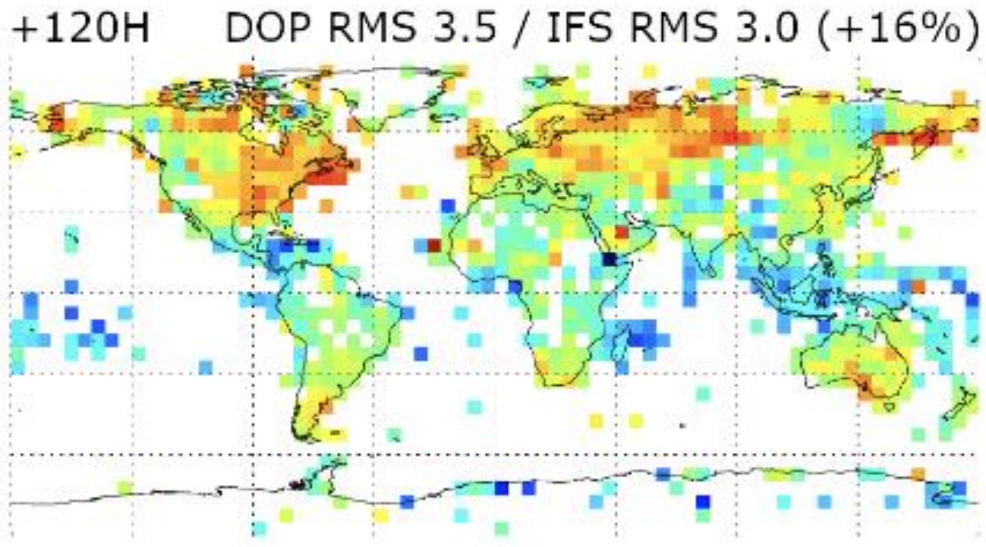

Normalized RMS difference in t2m forecast departures

AI-DOP better

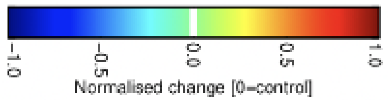

IFS 9km better

Comparing only obs processed by both systems!

Verification of gridded DOP forecasts (o96) **against reanalysis**

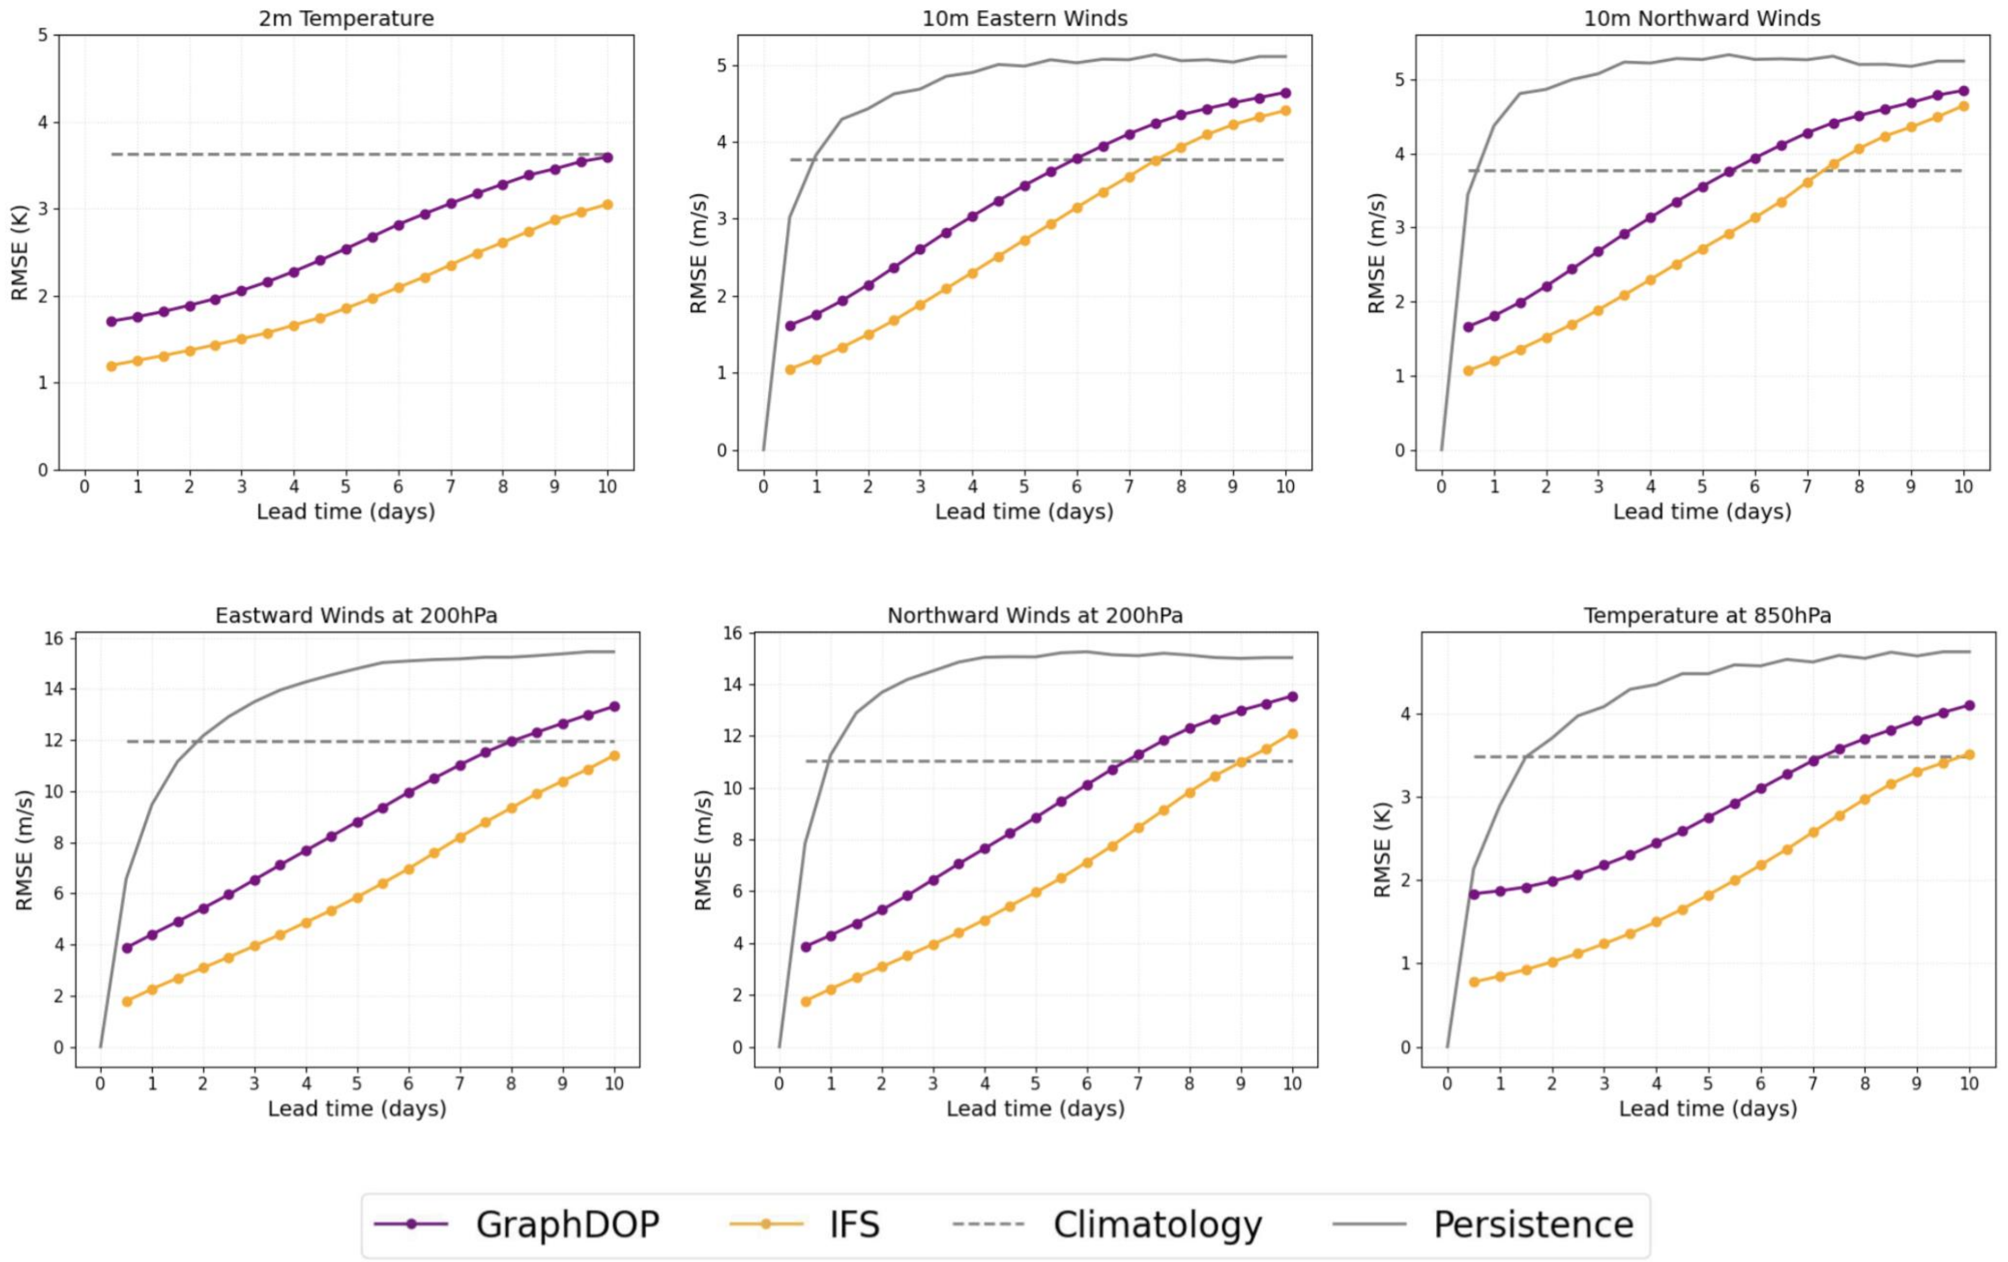

Use cases

A rapid freezing event in the Arctic (October 2022)

AMSR-2 channel 5 (10v)
brightness temperatures

12h forecast

10-day forecast

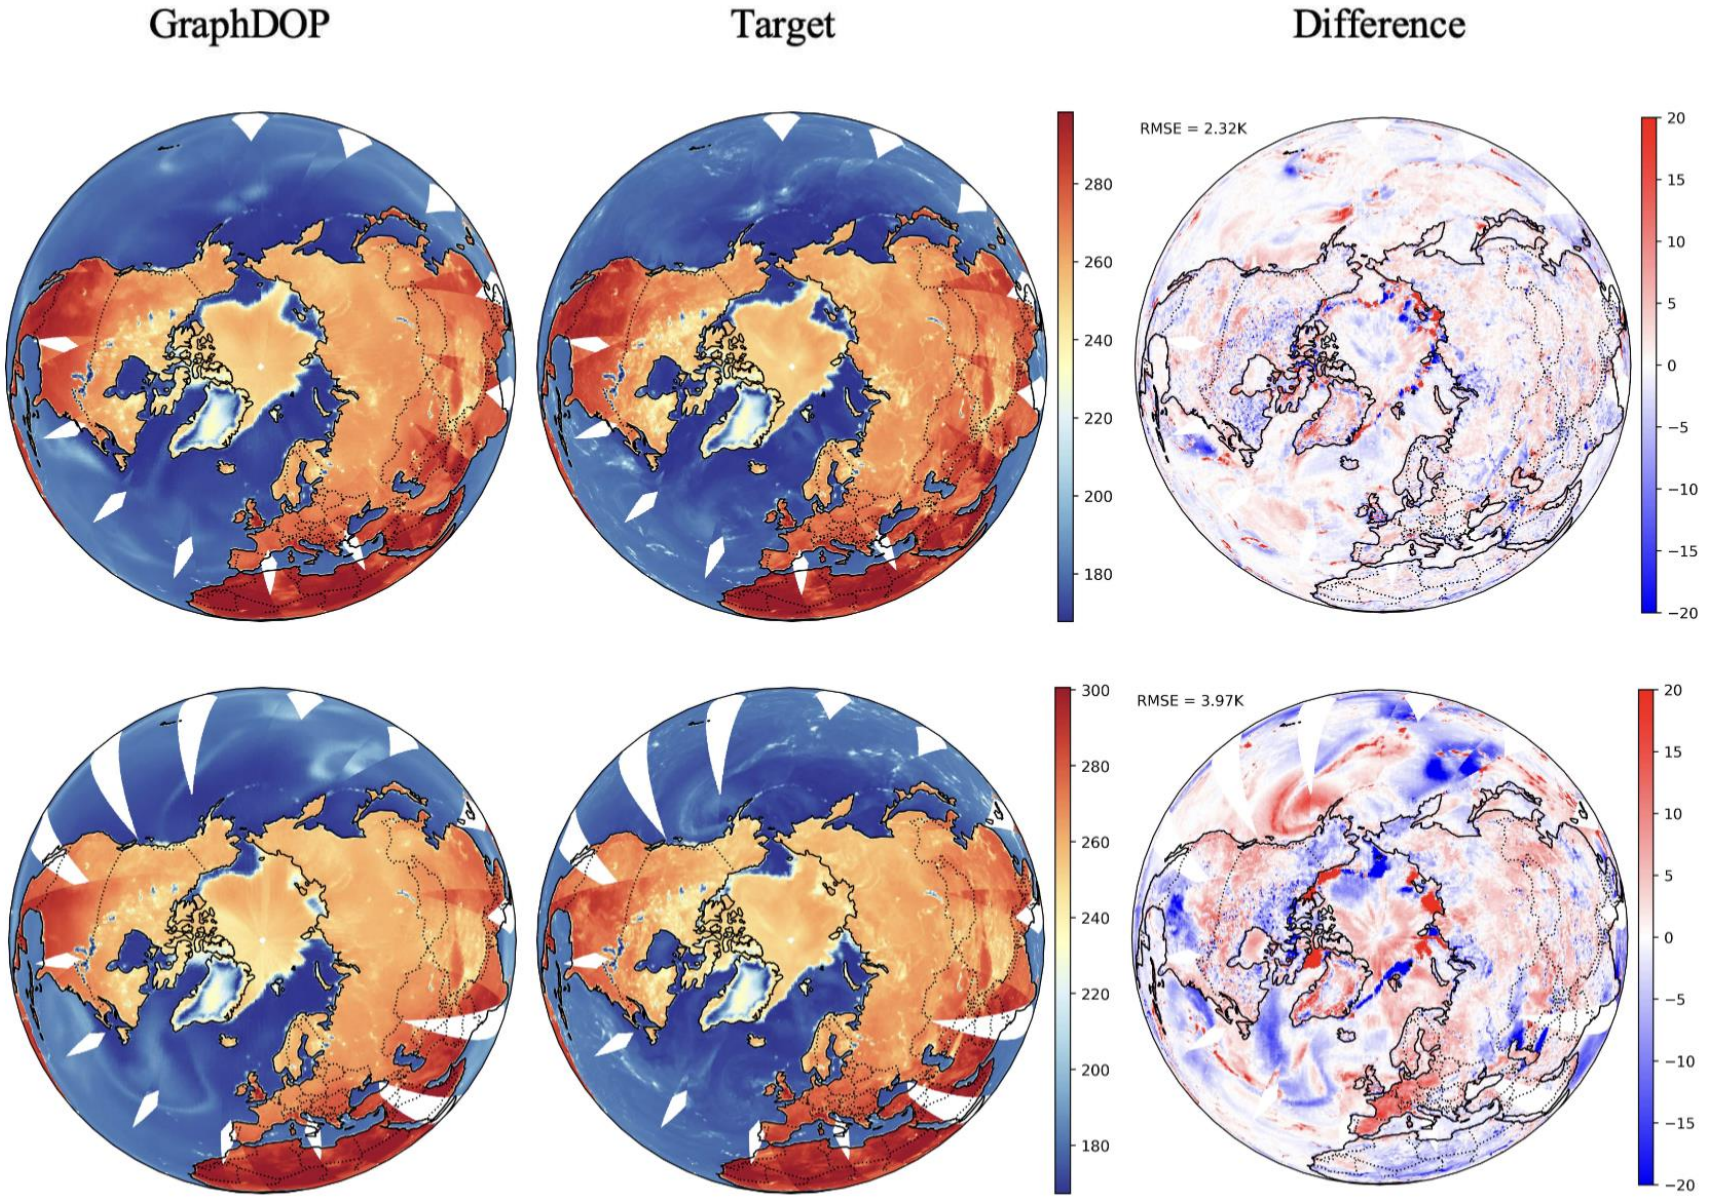

Hurricane Ian

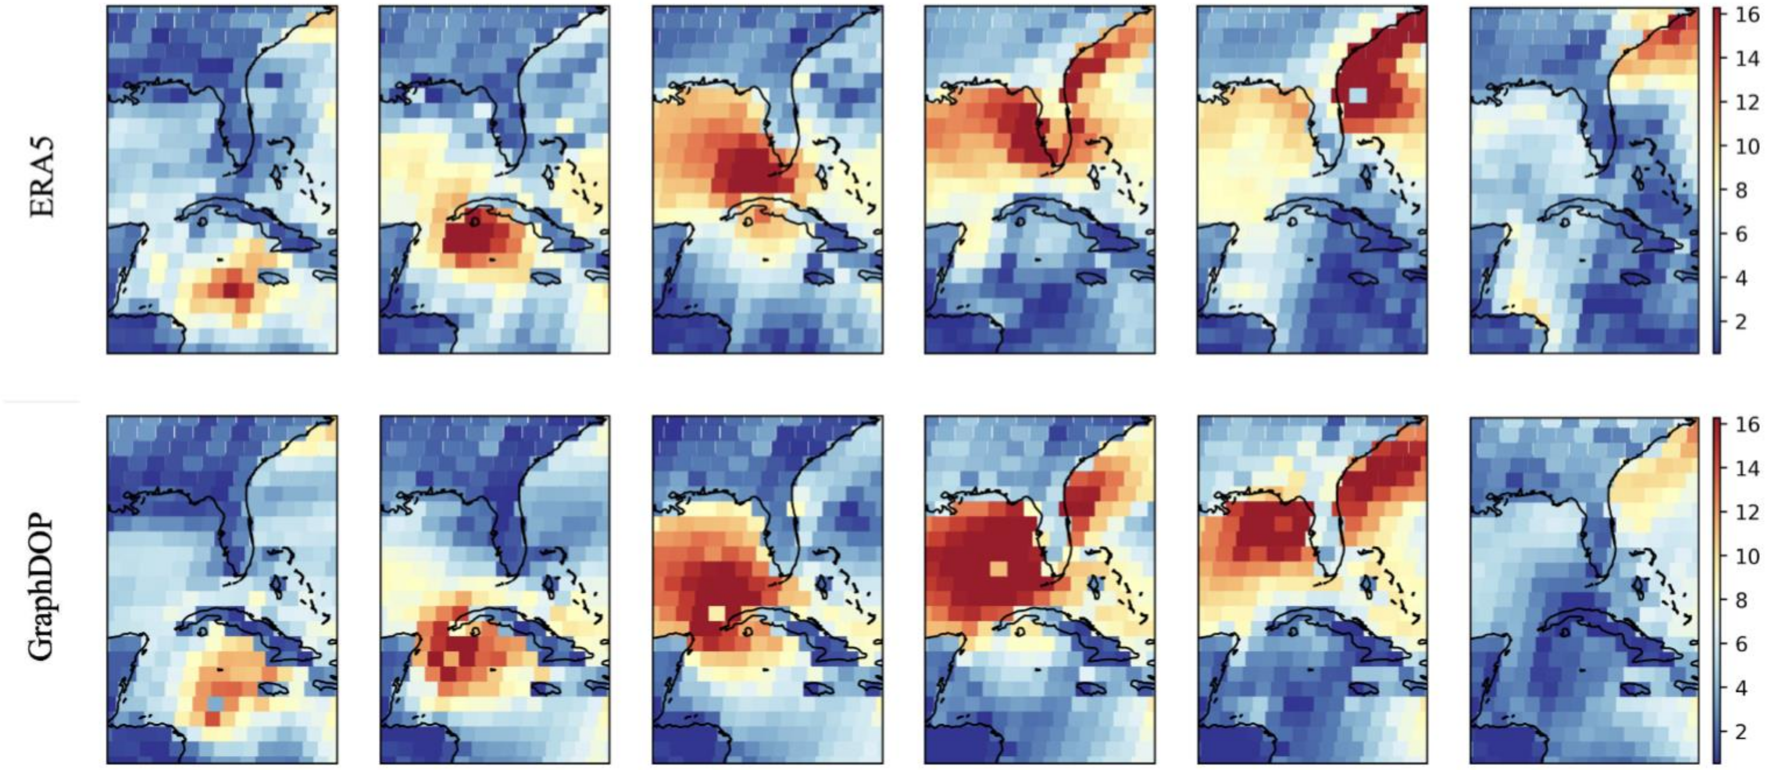

(b) Wind speed (m/s)

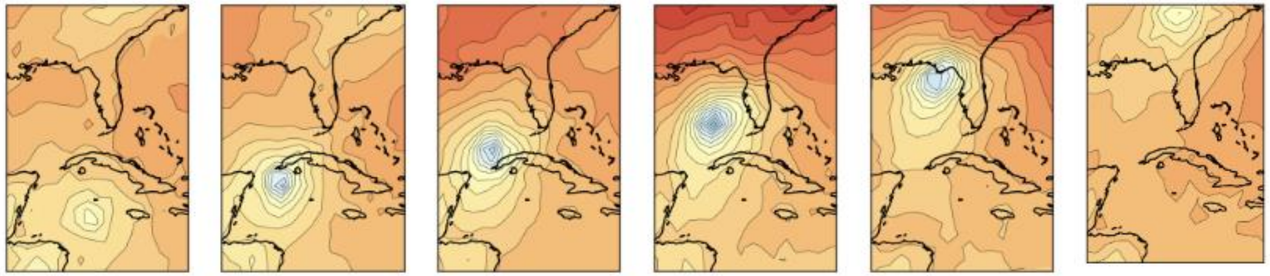

(a) Mean sea-level pressure (Pa)

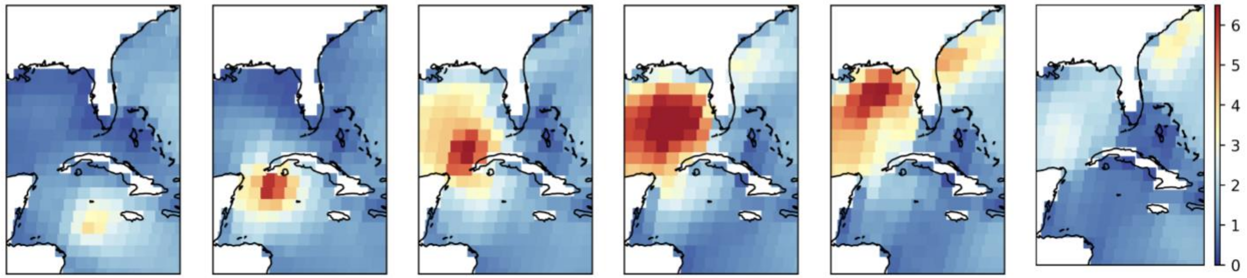

(c) Significant wave height (m)

Outlook

Ongoing work

Improving observation QC

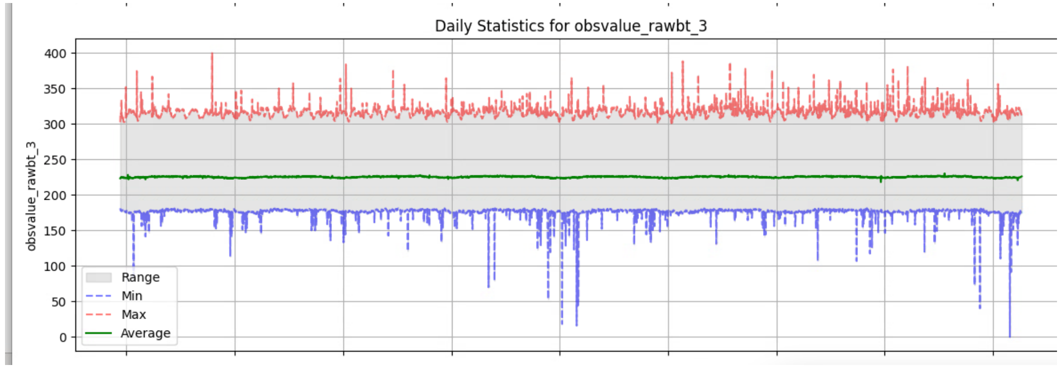

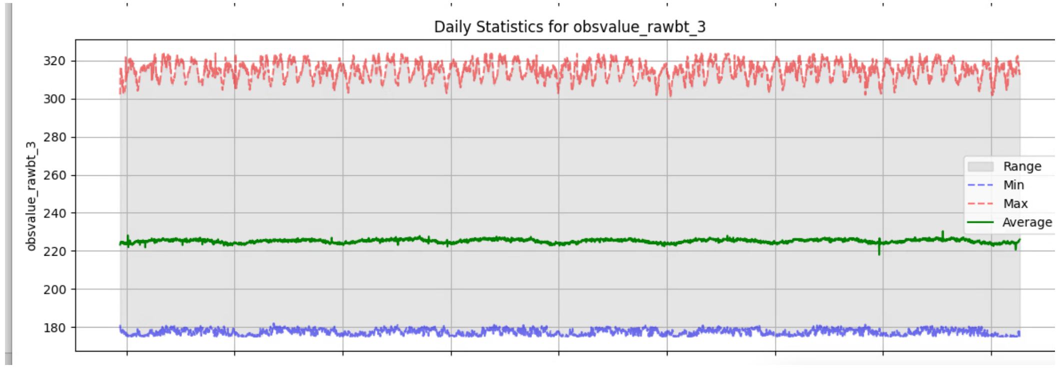

Probabilistic training (e.g., CRPS)

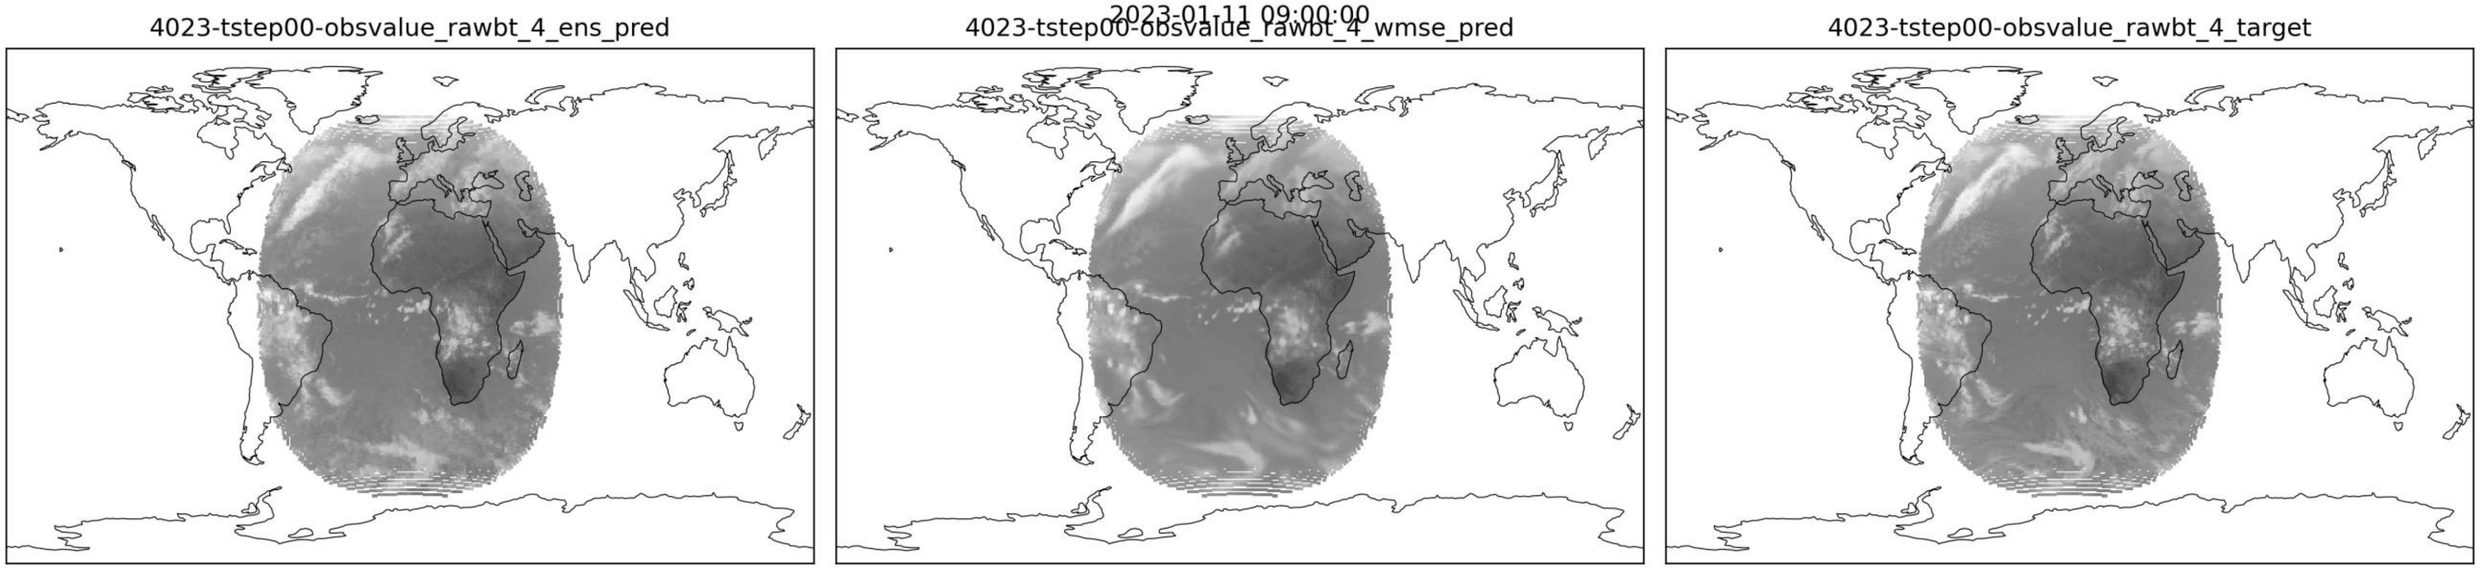

If AI-DOP works as well as we expect it to …

Bypass some of the most challenging aspects of conventional DA

Full Earth system forecasting without coupling, by learning simultaneously from atmosphere, ocean and land / cryosphere observations

Better exploitation of observations, even those with complex difficult physics…

The value of observations goes beyond initial conditions for physics-based NWP models – **they become the model**!

Sustainable R2O process… with the only dependency on observations!

If it doesn't …

Hybrid options (e.g., AIFS + observations) are already showing promise, also in a regional context

We will have learned a lot about the observations – this knowledge will carry over to hybrid designs

Other applications possible, e.g. real-time observation QC

Data driven weather forecasts trained and initialised directly from observations

A Preprint

Anthony McNally Christian Lessig Peter Lean Eulalie Boucher Mihai Alexe
Ewan Pinnington Matthew Chantry Simon Lang Chris Burrows Marcin Chrust
Florian Pinault Ethel Villeneuve Niels Bormann Sean Healy

European Centre for Medium-Range Weather Forecasts (ECMWF)

July 23, 2024

Abstract

Skilful Machine Learned (ML) weather forecasts have challenged conventional approaches to numerical weather prediction (NWP), demonstrating competitive performance compared to traditional physics-based approaches. Existing data-driven systems have been trained to forecast future weather by learning from long historical records of past weather, typically provided by reanalyses such as ECMWF's ERA5. These datasets have been made freely available to the wider research community, including the commercial sector, which has been a major factor in the rapid rise of ML forecast systems and the impressive levels of accuracy they have achieved. However, both historical reanalyses used for training and real-time analyses used for initial conditions are produced by data assimilation, essentially an optimal blending of observations with a traditional physics-based forecast model. As such, many ML forecast systems have an implicit, unknown and unquantified dependence on the physics-based models they seek to challenge. Here we propose a new and radical approach to weather forecasting, by training a neural network to predict future weather purely from historical observations with no dependence on a physics-based model or reanalysis datasets. We use raw observations (level-1) to initialise a model of the atmosphere

GraphDOP: Towards skilful data-driven medium-range weather forecasts learnt and initialised directly from observations

A Preprint

Mihai Alexe Eulalie Boucher Peter Lean Ewan Pinnington Patrick Laloyaux
Anthony McNally Simon Lang Matthew Chantry Chris Burrows Marcin Chrust
Florian Pinault Ethel Villeneuve Niels Bormann Sean Healy

European Centre for Medium-Range Weather Forecasts (ECMWF)

December 20, 2024

Abstract

We introduce GraphDOP, a new data-driven, end-to-end forecast system developed at the European Centre for Medium-Range Weather Forecasts (ECMWF) that is trained and initialised exclusively from Earth System observations, with no physics-based (re)analysis inputs or feedbacks. GraphDOP learns the correlations between observed quantities - such as brightness temperatures from polar orbiters and geostationary satellites - and geophysical quantities of interest (that are measured by conventional observations), to form a coherent latent representation of Earth System state dynamics and physical processes, and is capable of producing skilful predictions of relevant weather parameters up to five days into the future.

1 Introduction

In recent years, data-driven approaches to numerical weather prediction (NWP) have taken the field by storm, with several global models demonstrating forecast skill scores comparable or superior to that of leading physics-based NWP systems across a wide range of weather variables and lead times [Pathak et al., 2022, Lam et al., 2023, Bi et al., 2023, Bodnar et al., 2024, Lang et al., 2024a]. Without exception, these data-driven models have been trained on reanalysis products such as ECMWF's ERA5 [Hersbach et al., 2020]. To produce a forecast, the models must be started from a weather (re)analysis valid at the initial time of the forecast.

A (re)analysis is the product of data assimilation, a family of algorithms that aim to optimally combine the best available estimate of the current global atmospheric state - e.g., a previous short-range forecast from a physics-based weather model - with information obtained from Earth System observations. For example, the ECMWF runs four-dimensional