Building Machine Learning Limited Area Models

Kilometer-Scale Weather Forecasting in Realistic Settings

Joel Oskarsson
Linköping University

Simon Adamov
MeteoSwiss, ETH Zurich

LINKÖPING UNIVERSITY MeteoSwiss ETH zürich

Machine Learning Weather Prediction

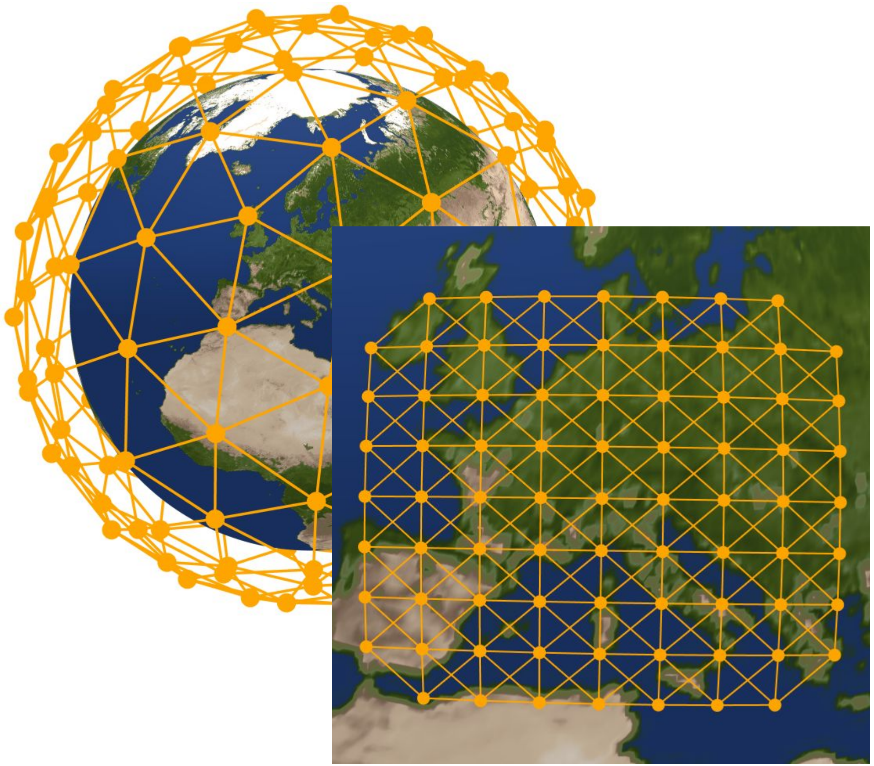

Regional Models:

- Higher resolution
- Model challenging processes
- Utilize regional data

The Limited Area Forecasting Problem

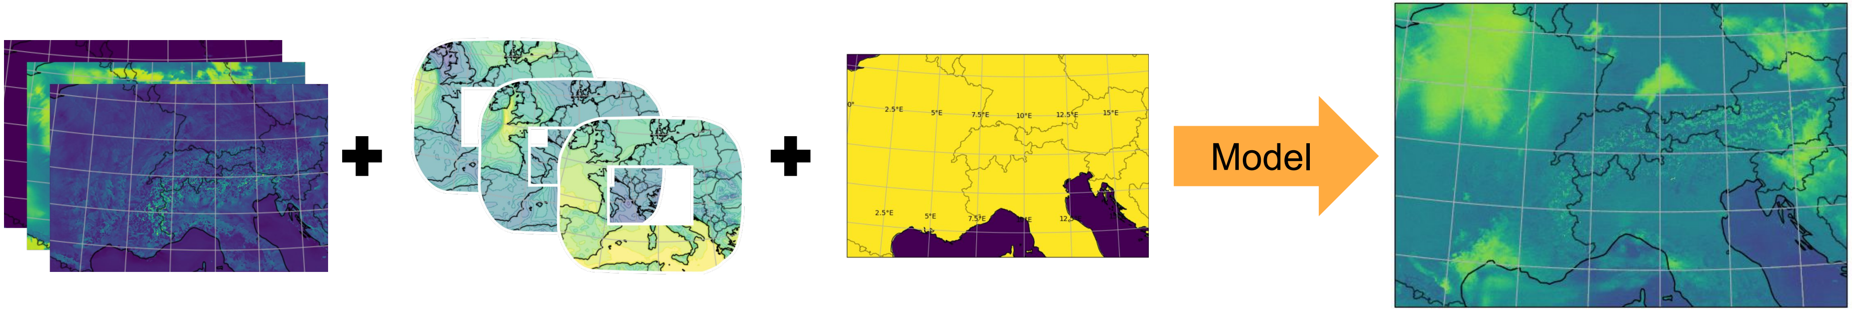

Initial conditions
Regional

Boundary forcing
From global model

Forcing + static features

Forecast
Regional

This project: Realistic setting, forced by global forecast (IFS)

A Framework for Machine Learning LAMs

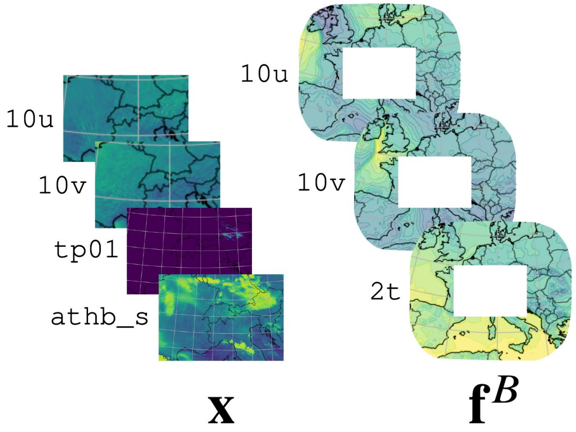

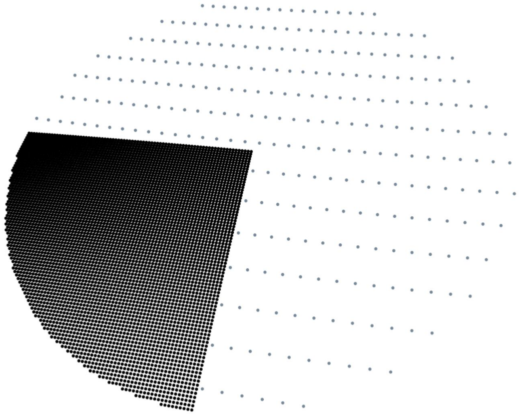

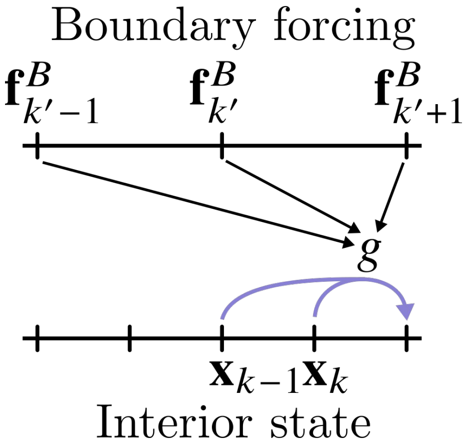

Different atmospheric variables

Different grid layouts

Different length of time steps

Graph-based Model

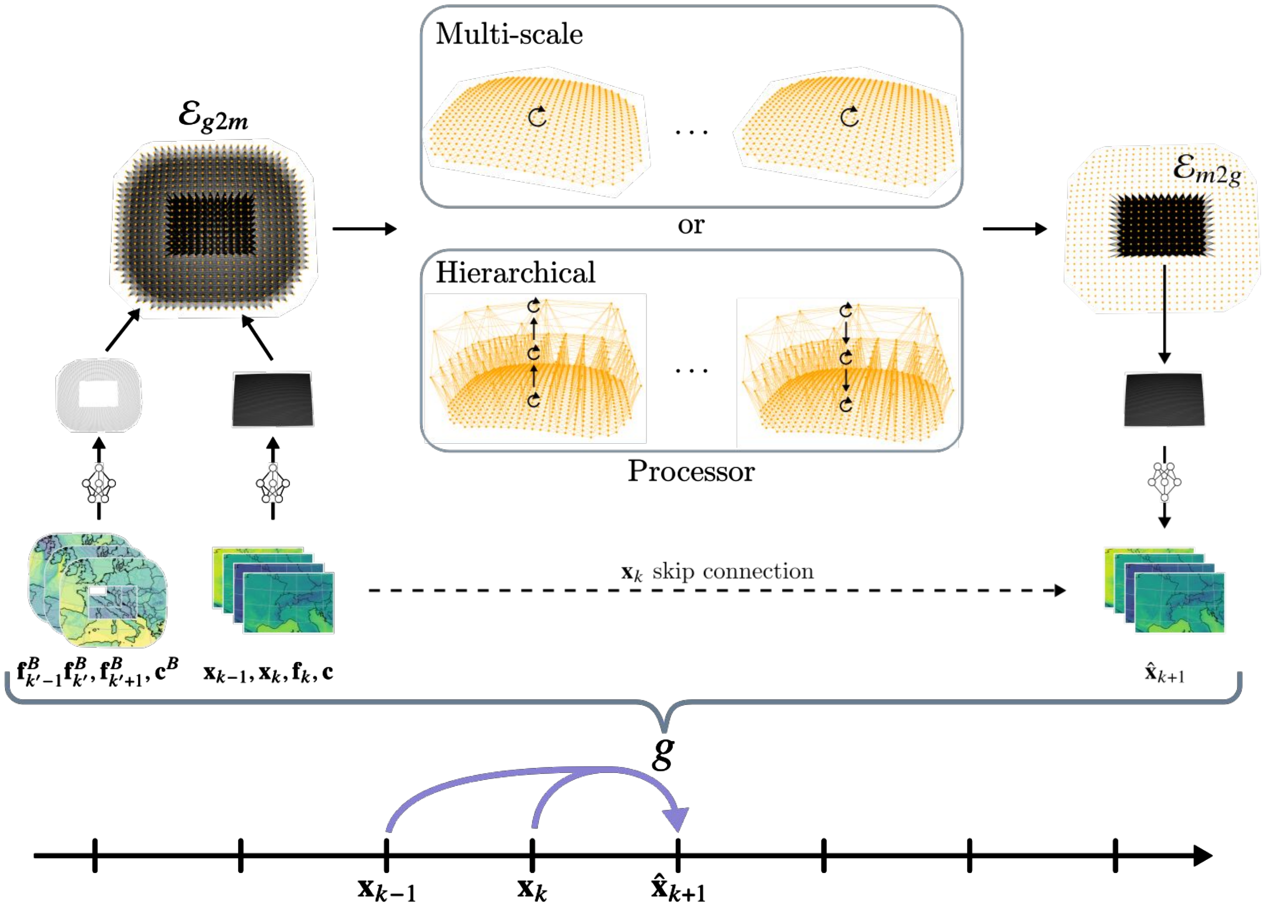

Datasets

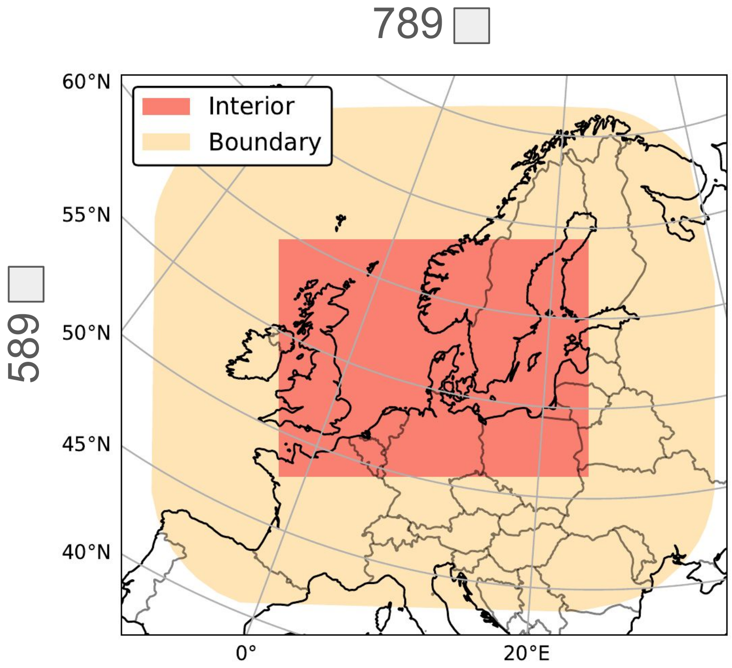

DANRA - 2.5km - 3h - 21 years

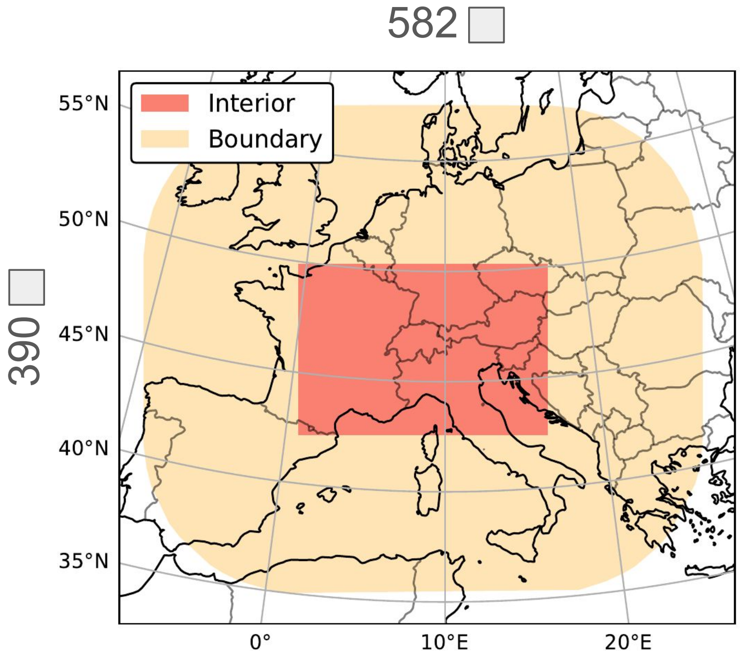

COSMO - 2.2km - 1h - 4 years

Boundary forcing: ERA5/IFS - 0.25° - 6h

Experiments

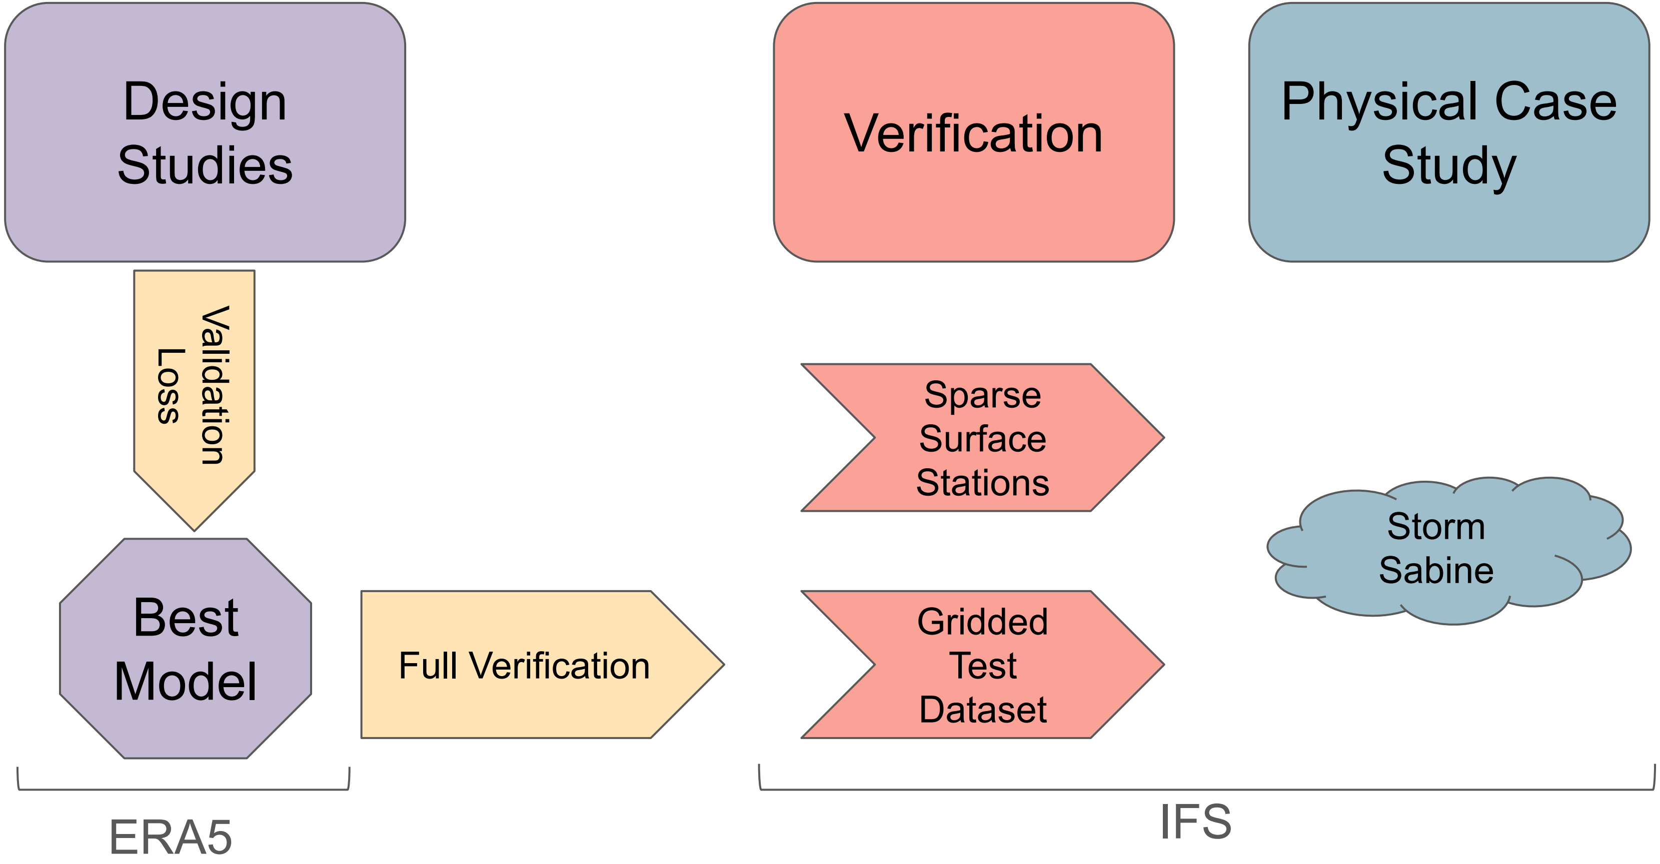

Design Studies - Boundary Type

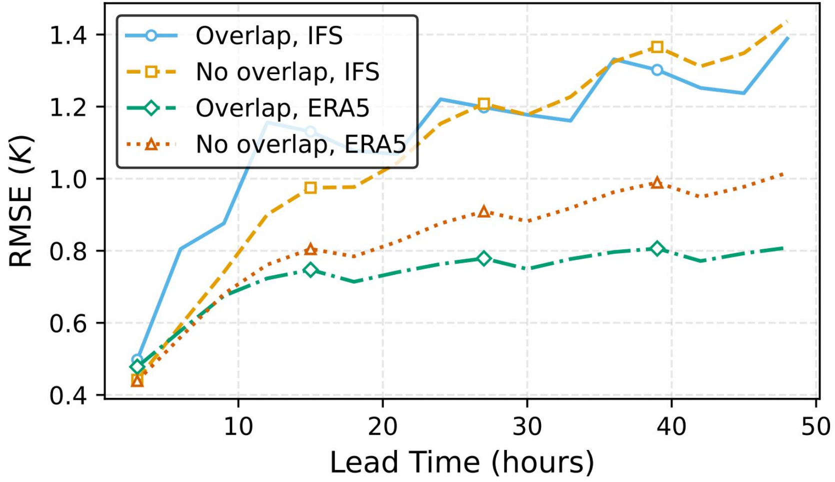

(a) 2 m temperature (2t)

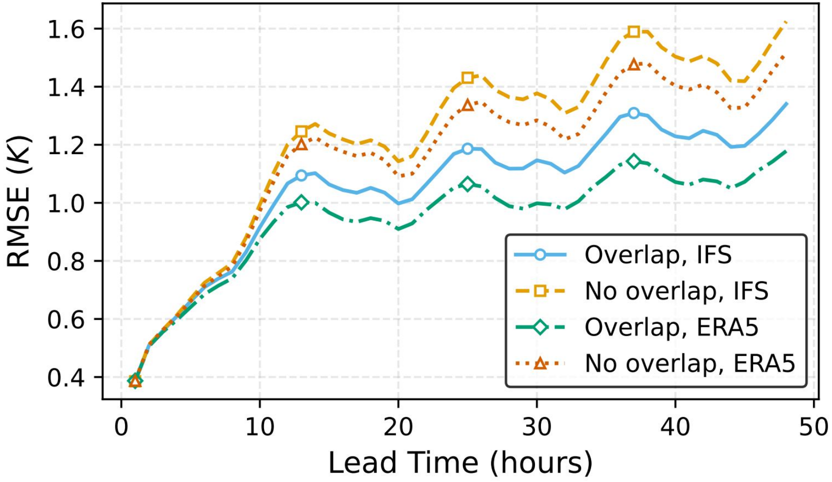

(a) 2 m temperature (2t)

DANRA
Overlap less important
Growing error inherited from IFS

COSMO
Overlap more important
Forecasting vs. downscaling

Design Studies - Boundary Width - COSMO

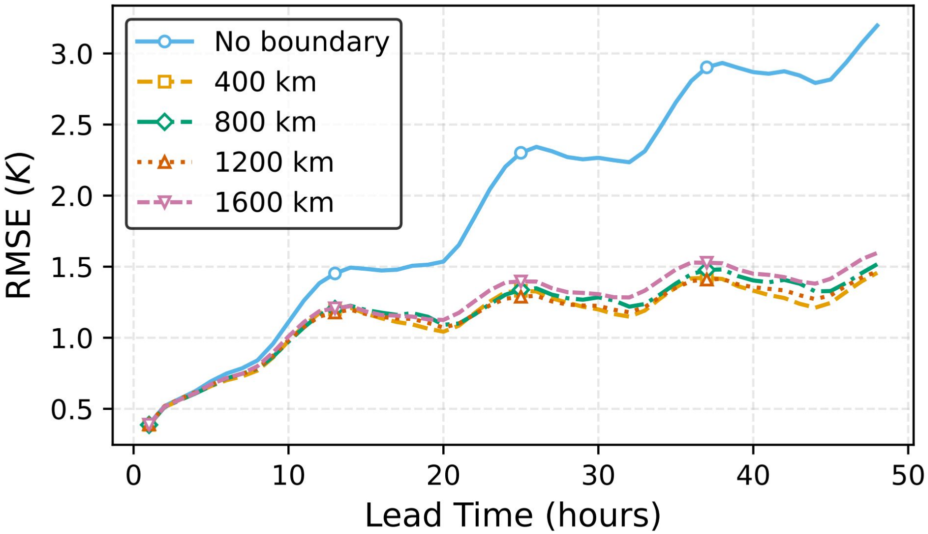

(a) 2 m temperature (2t)

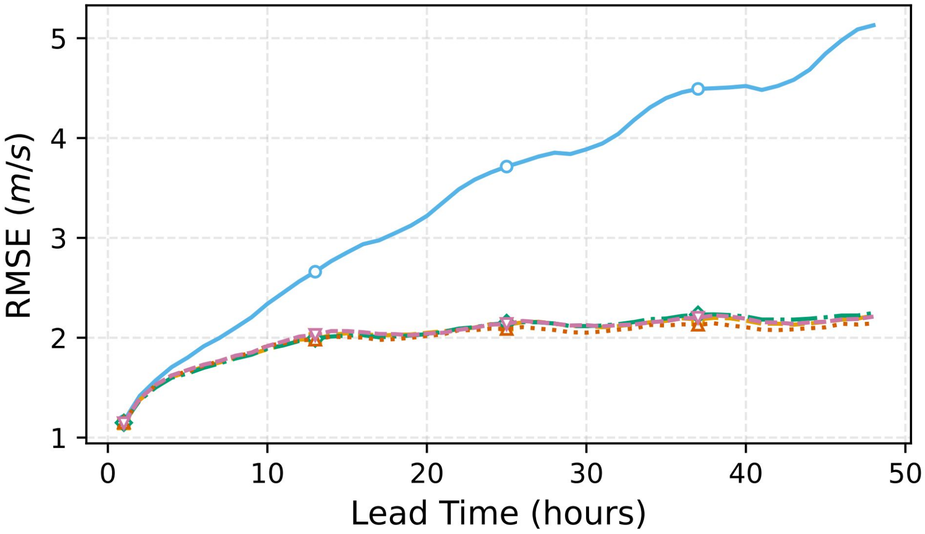

(b) 10 m wind

The boundary is important but can be smaller than we expected

Verification Sparse - COSMO

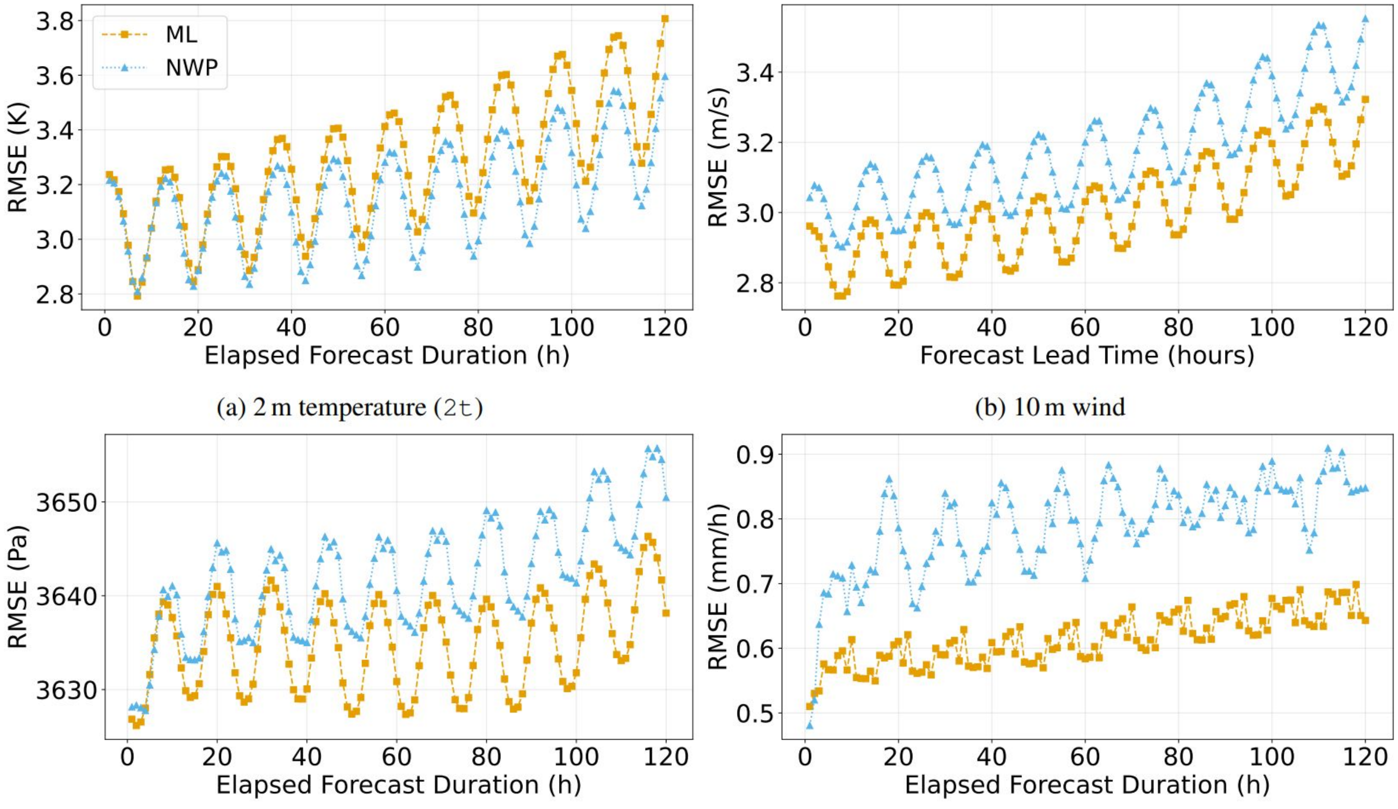

(a) 2 m temperature (2t)

(b) 10 m wind

(c) Surface pressure (sp)

(d) Precipitation (tp01)

Verification Sparse - Threshold Based - COSMO

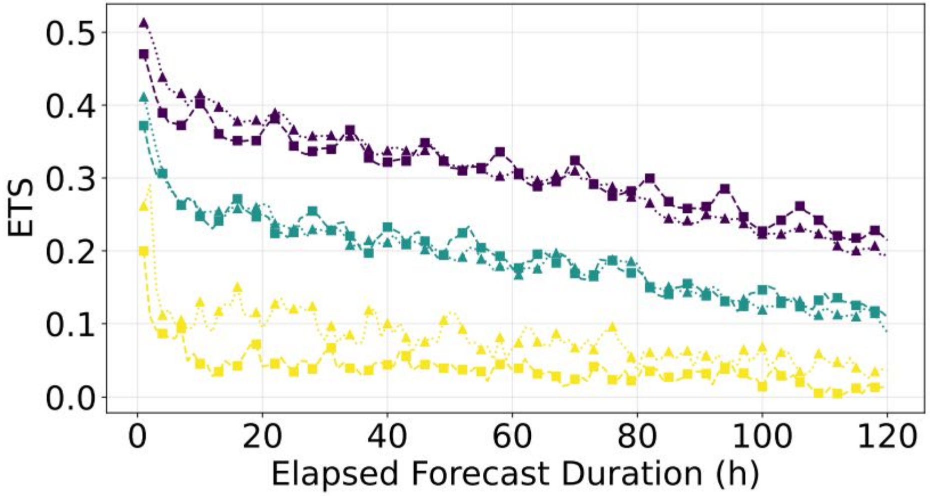

(d) Precipitation (`tp01`) ETS

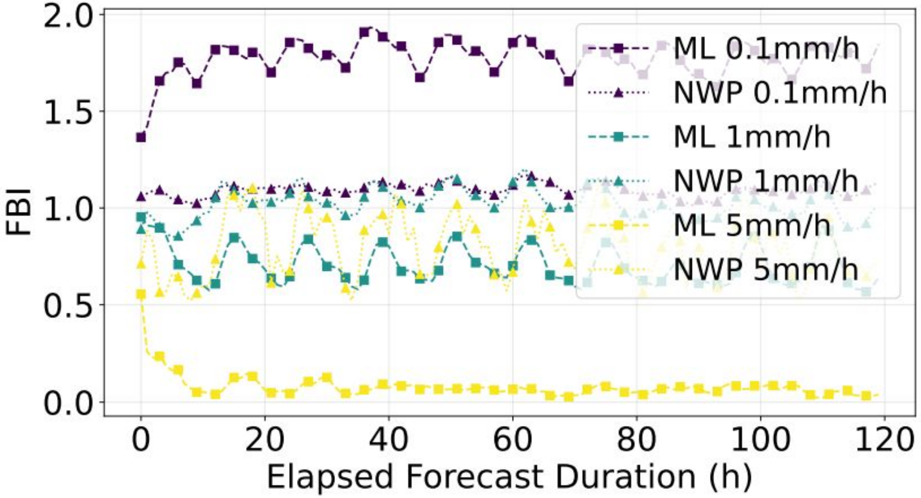

(c) Precipitation (`tp01`) FBI

Verifying the model with appropriate figures and metrics is crucial

Verification Gridded - Energy Spectra - DANRA

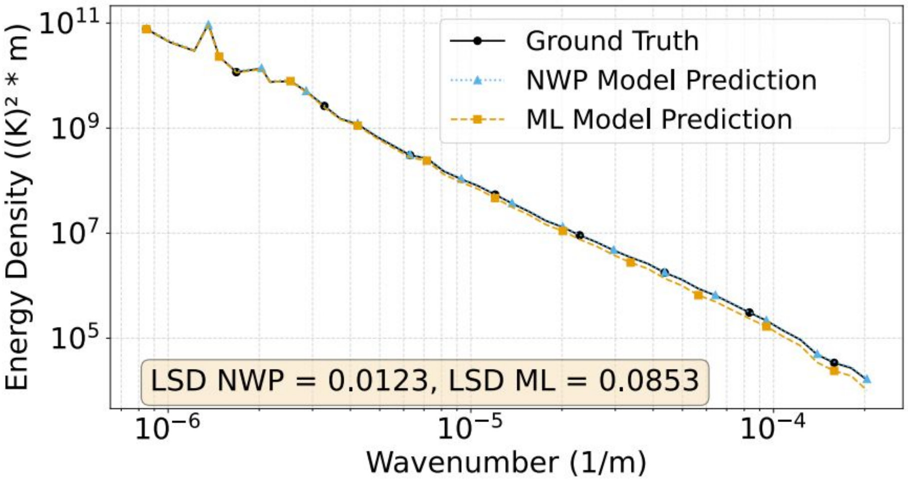

(a) 2 m temperature (2t)

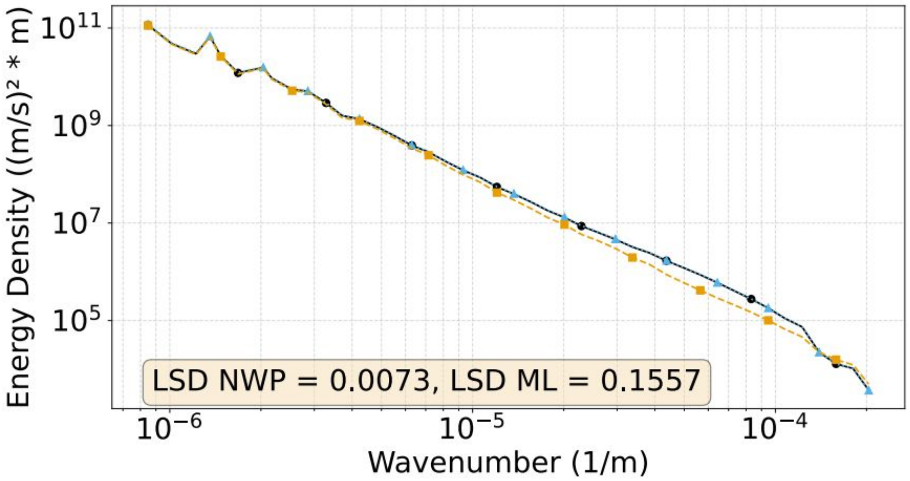

(b) 10 m wind-component (10u)

Energy spectra showing only slight over-smoothing of high frequencies

Case Study Sabine - COSMO

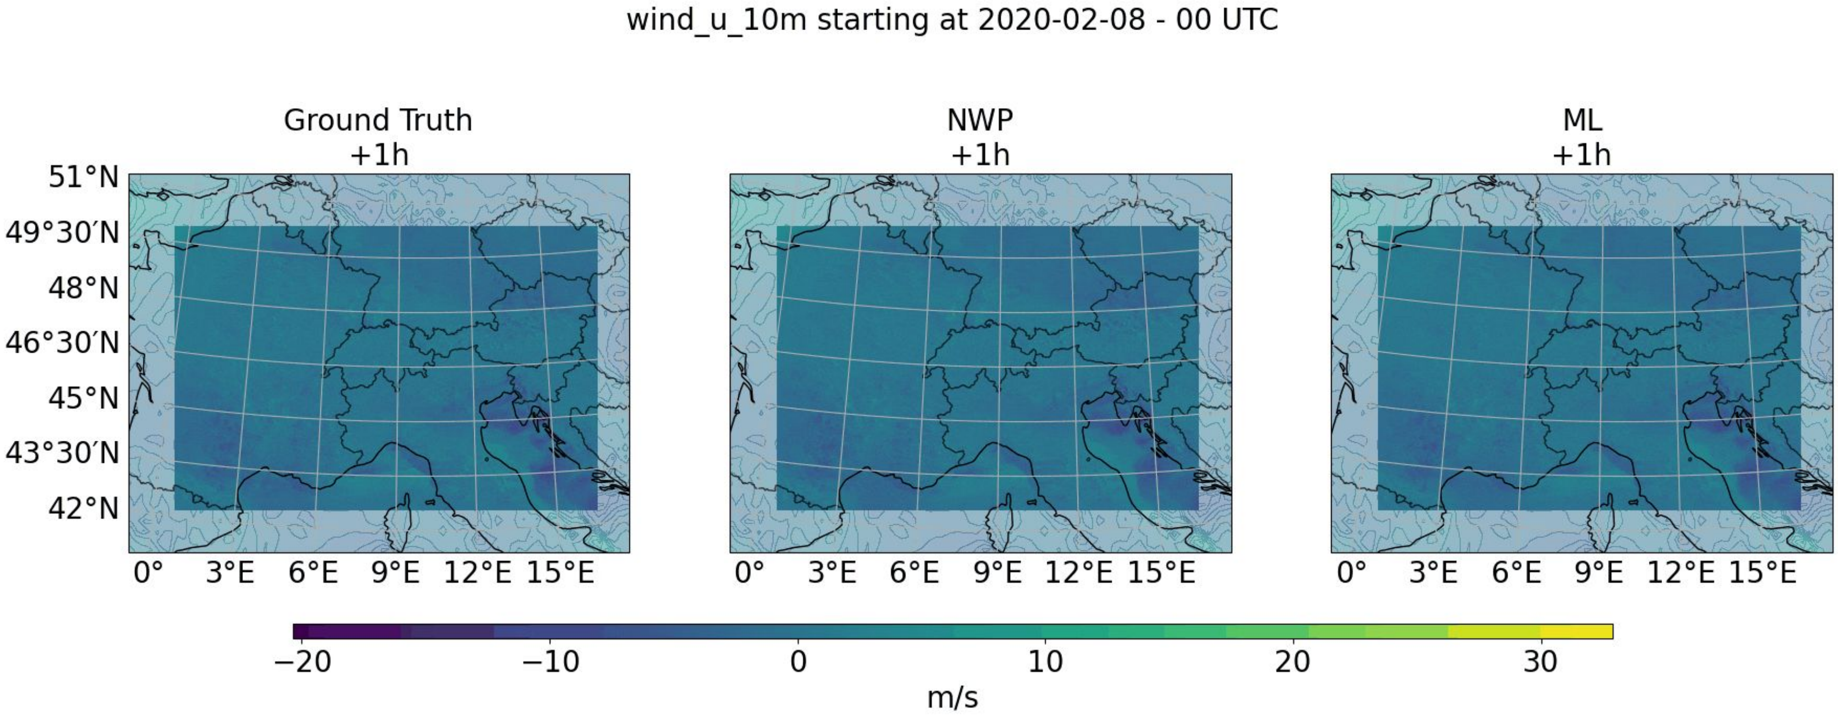

Case Study Sabine - DANRA

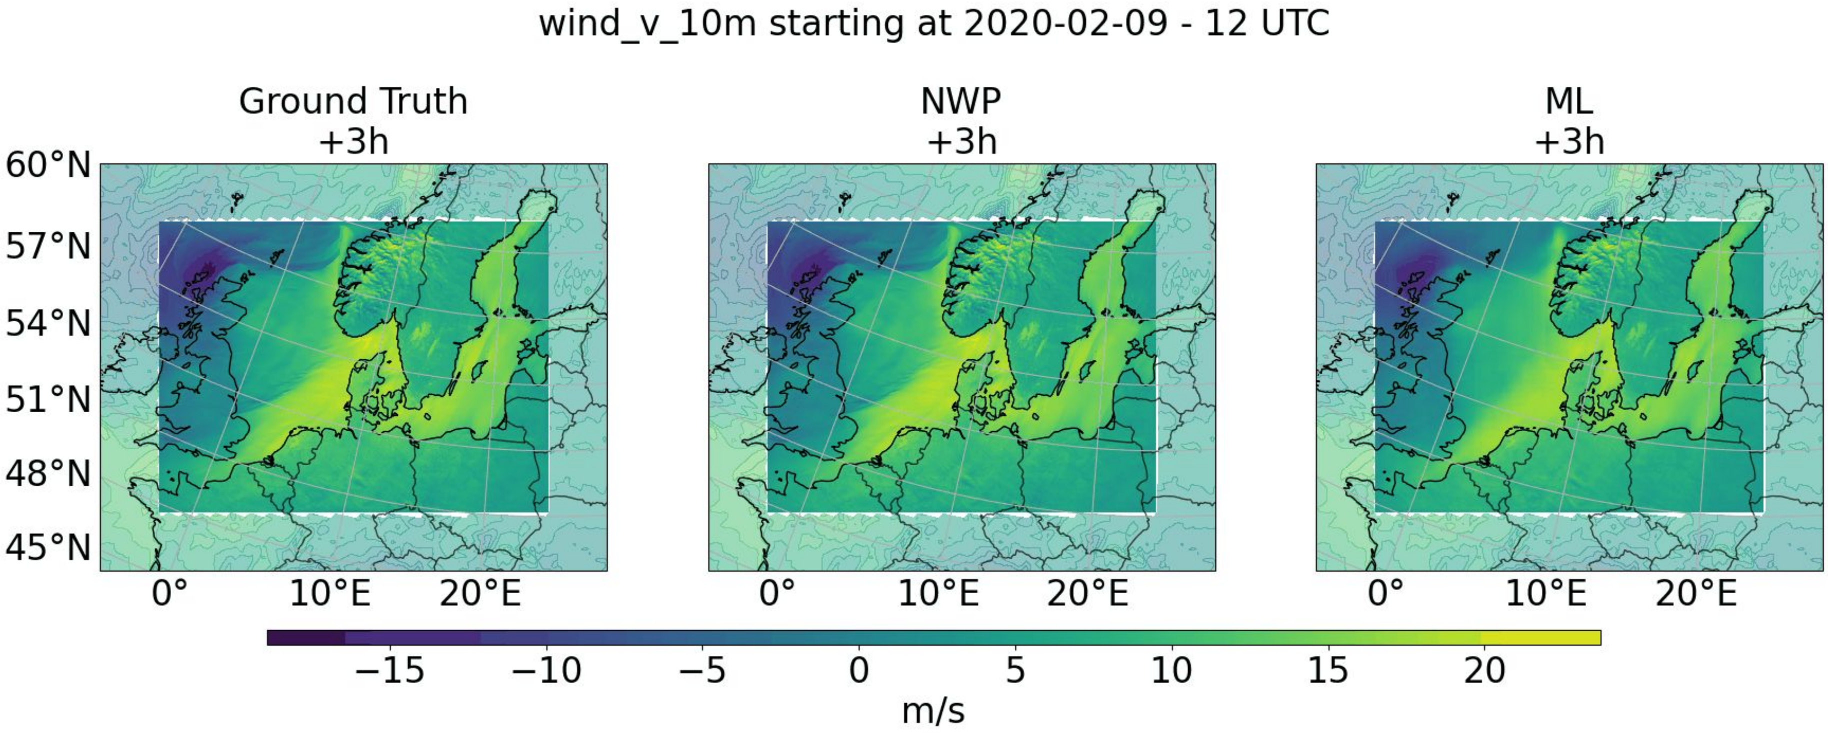

Outlook

Probabilistic LAMs[1]

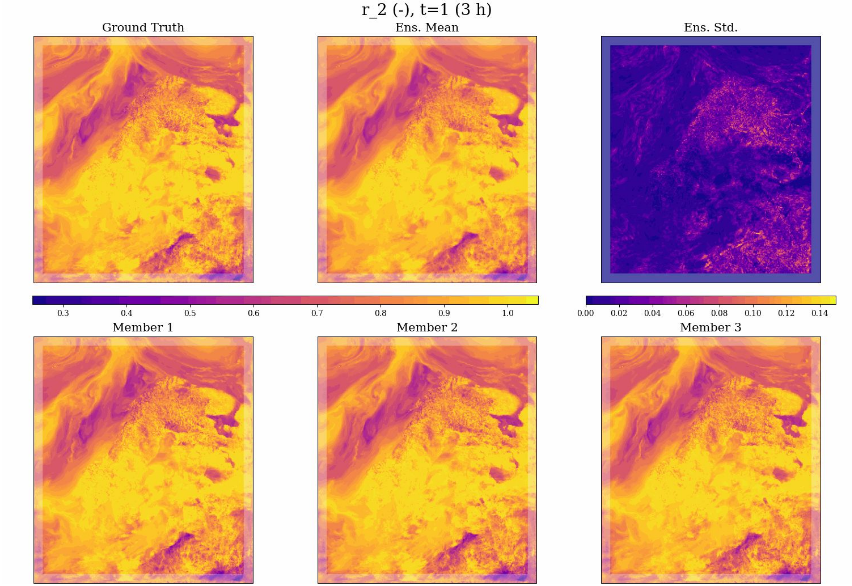

Regional Earth System Models

- Oceanography[2]

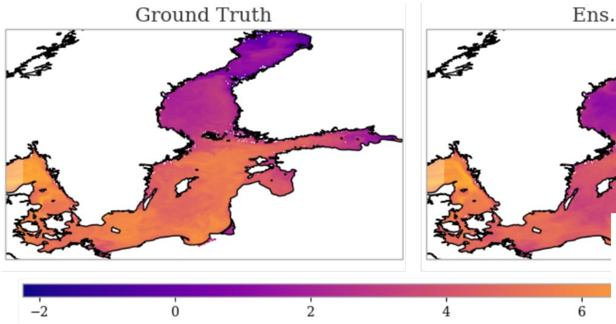

Temperature, 0.5 m (C)

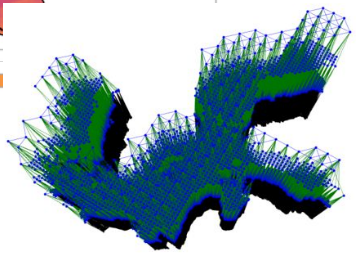

[2] D. Holmberg, et al. (2025). Accurate Mediterranean Sea forecasting via graph-based deep learning. Preprint. + Ongoing work.

[1] J. Oskarsson, et al. (2024). Probabilistic Weather Forecasting with Hierarchical Graph Neural Networks. NeurIPS.
E. Larsson, et al. (2025). Diffusion-LAM: Probabilistic Limited Area Weather Forecasting with Diffusion. CCAI Workshop @ ICLR.
J. Pathak, et al. (2024). Kilometer-scale convection allowing model emulation using generative diffusion modeling. Preprint.

Building Machine Learning Limited Area Models

Kilometer-Scale Weather Forecasting in Realistic Settings

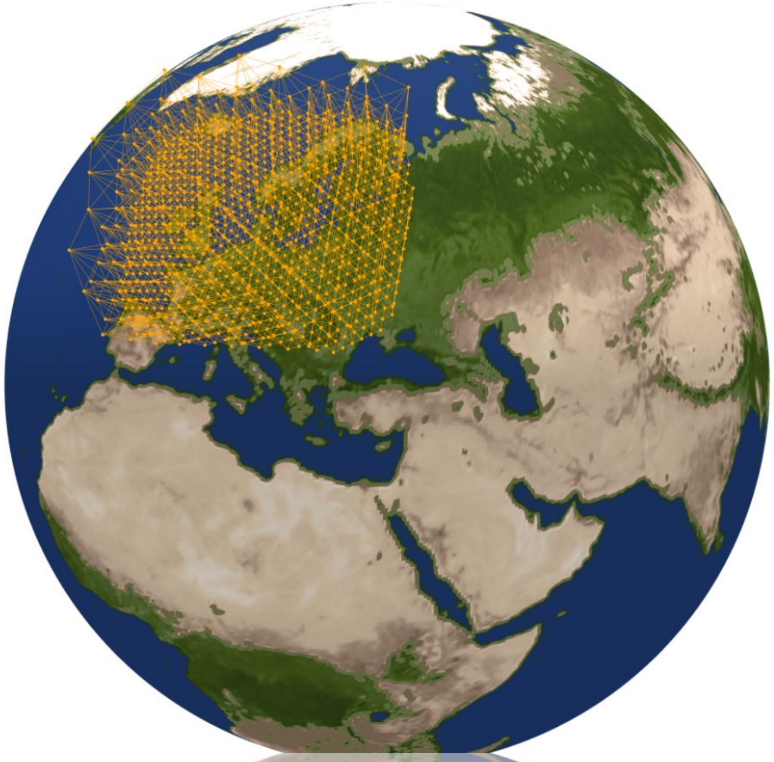

Simon Adamov (MeteoSwiss, ETH)
Joel Oskarsson (LiU)

Co-authors:
Leif Denby (DMI)
Tomas Landelius (SMHI, LiU)
Kasper Hintz (DMI)
Simon Christiansen (DMI)
Irene Schicker (GeoSphere)
Carlos Osuna (MeteoSwiss)
Fredrik Lindsten (LiU)
Oliver Fuhrer (MeteoSwiss)
Sebastian Schemm (ETH, Cambridge)

Paper: https://arxiv.org/abs/2504.09340
Code: https://github.com/joeloskarsson/neural-lam-dev/releases/tag/building-ml-lams

Design Studies - Graph Design - DANRA

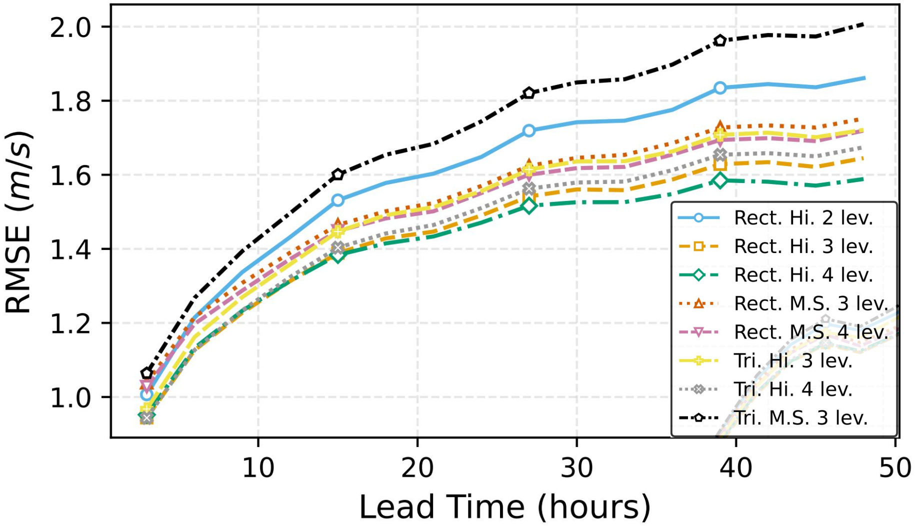

(b) 10 m wind

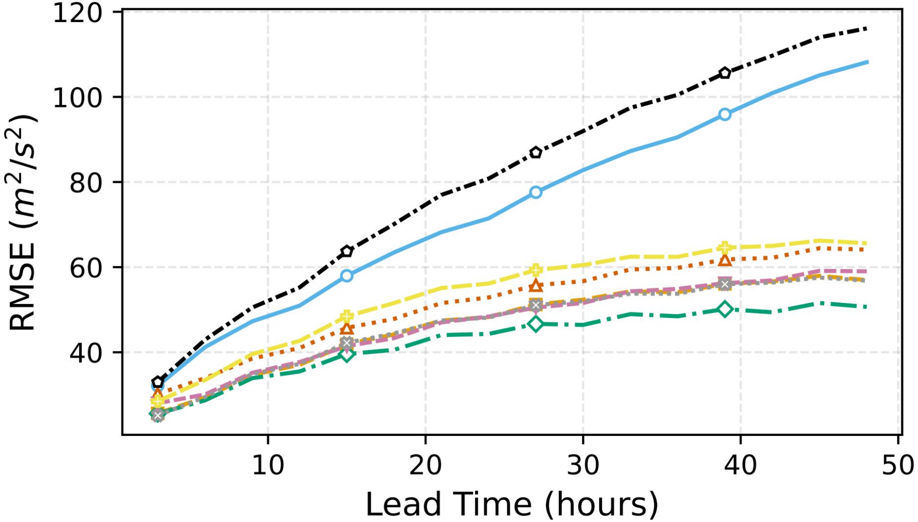

(d) Geopotential at 600 hPa (z600)

4-level hierarchical rectangular graph slightly outperforms the others

Verification Gridded - Precipitation - COSMO

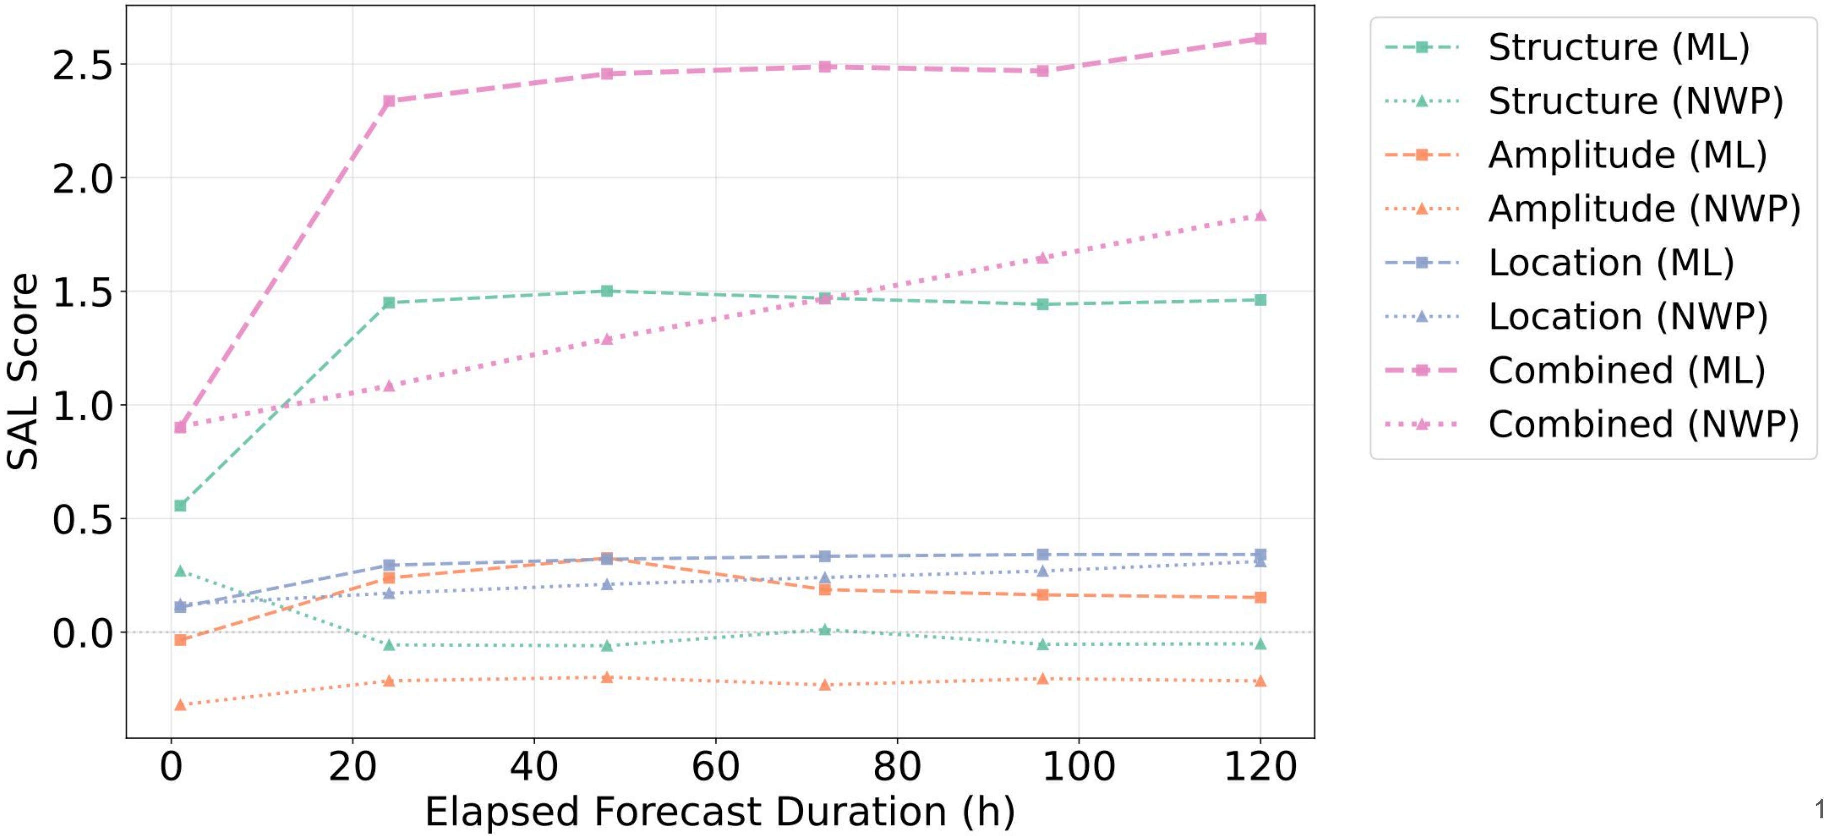

Verification Gridded - Vertical Profiles

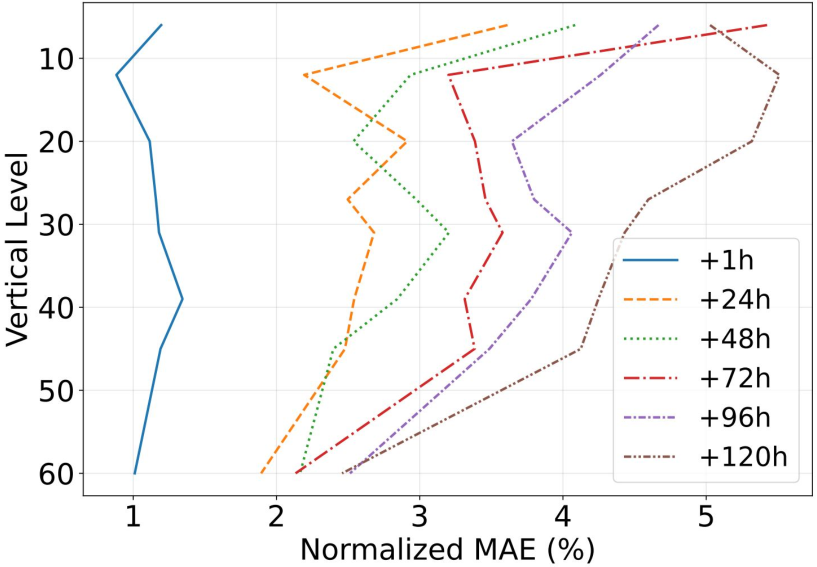

Win u-component

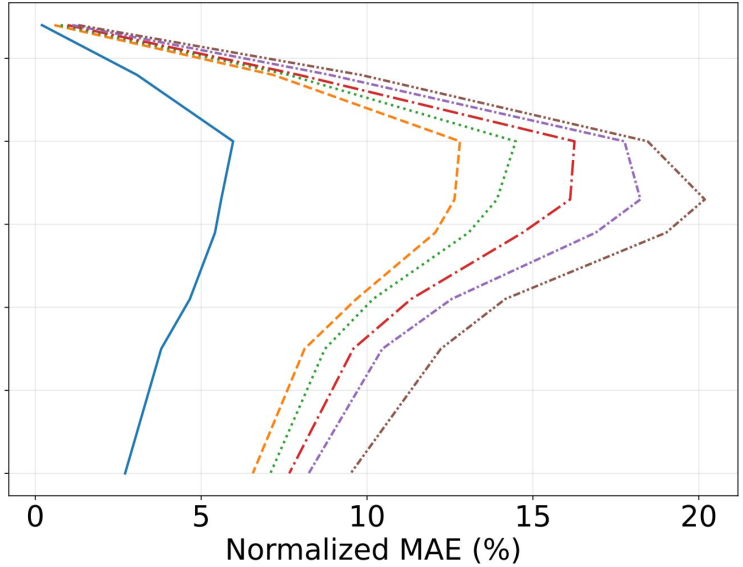

Relative Humidity

Shown are forecasts for COSMO - DANRA is more balanced

- Model performs worse in the upper atmosphere
- Loss weights can be adjusted

HPC Resources

Configuration	DANRA	COSMO
Dataset size (time steps)	54,896	33,660
GPU Type	NVIDIA A100/H100	NVIDIA GH200
Total GPUs per run	16	256 (H100s)
Pre-training Phase		
Epochs	80	200
Autoregressive rollout steps	1	1
Average training time [h]	144	12
Total GPU hours	2,304	3,072
Fine-tuning Phase		
Epochs	3	50
Autoregressive rollout steps	4	4^*
Average training time [h]	36	14
Total GPU hours	576	3,584
Number of Trainings		
Pre-training	12	9
Fine-tuning	12	12
Total GPU-hours	34,560	70,656

Graph design

Rectangular | Triangular

Multi-scale

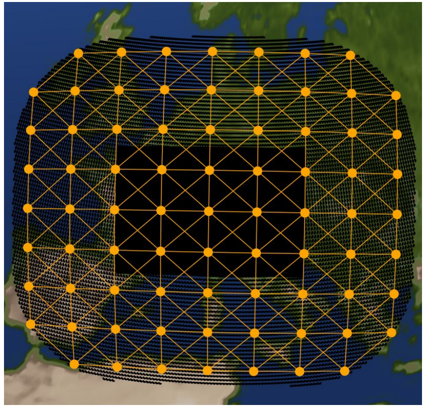

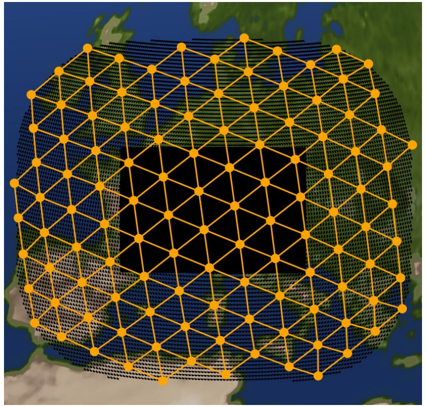

Hierarchical