Earth Sciences Department

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

eMONARCH: a deep learning emission-sensitive chemical-transport model

EGU General Assembly

04/05/2026

Michael Orieux, David Mathas, Isidre Mas Magre, Hervé Petetin

Context and motivation

MONARCH: state-of-the-art physics-based chemistry-transport model (CTM), **used for reseach, operational and policy-oriented studies**

To address the computational cost of CTMs like MONARCH, we want to develop **deep-learning-based surrogate models** (orders of magnitude faster while accurate enough to be useful).

Main foreseen applications:

- **Characterisation of emission sectors contribution**
- **Assessment of emission scenarios impact in air quality planning**
- **Quantification of CTM uncertainties for data assimilation**

Today we present preliminary results on predicting **only NO2 and NO concentrations**

MONARCH CTM
(expensive, accurate)

learn to emulate →

Deep Neural Network
(fast, approximate)

Inputs and targets, and training design

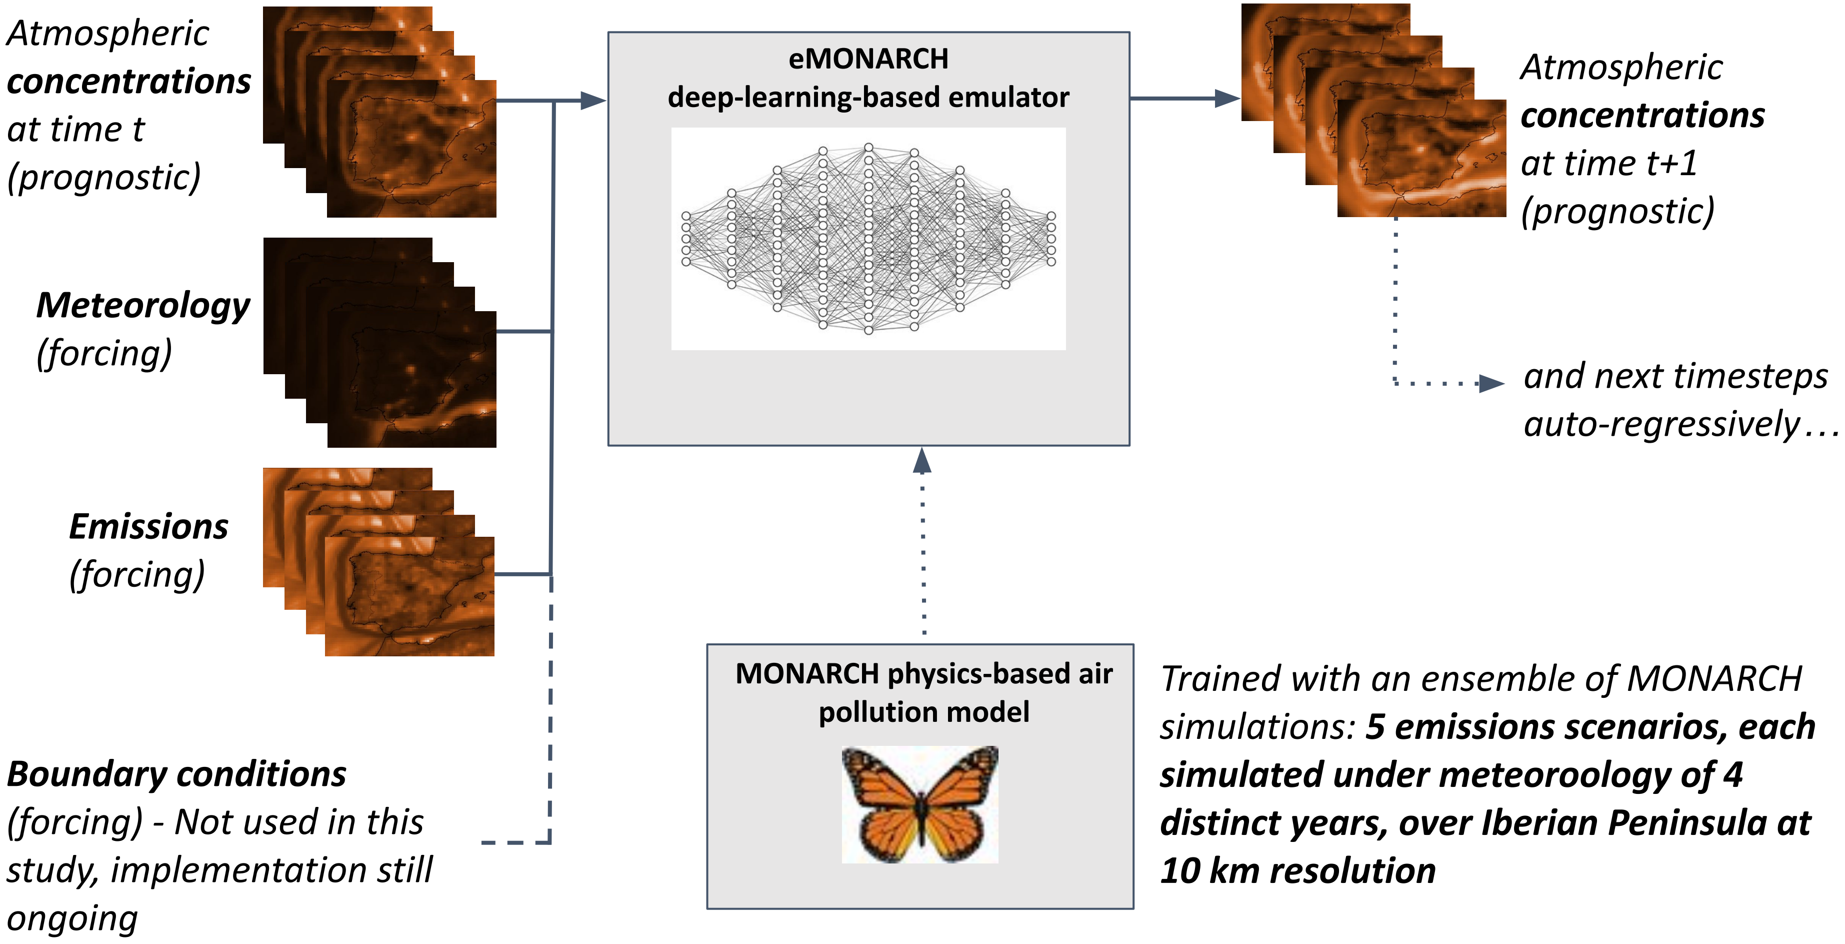

Architectures

First architecture: UNet

Train with **RMSE** loss and **cosine** learning rate scheduler for ~ 900 epochs with AdamW

In the encoder - decoder CNN architecture, we tuned the **base number of filters**

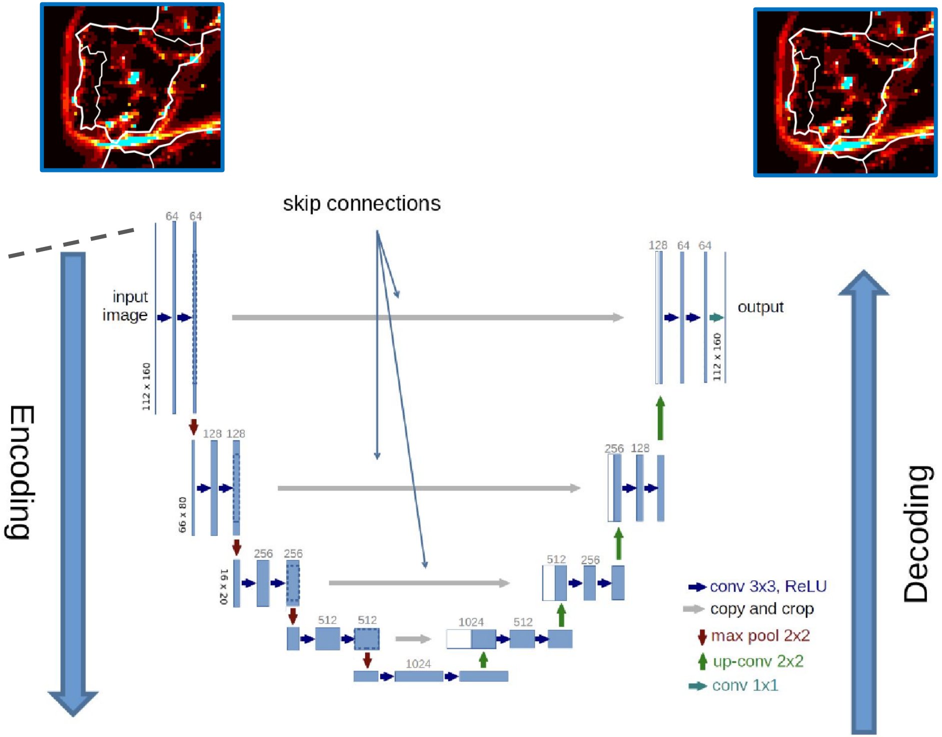

Second architecture: Graph Neural Network (GNN)

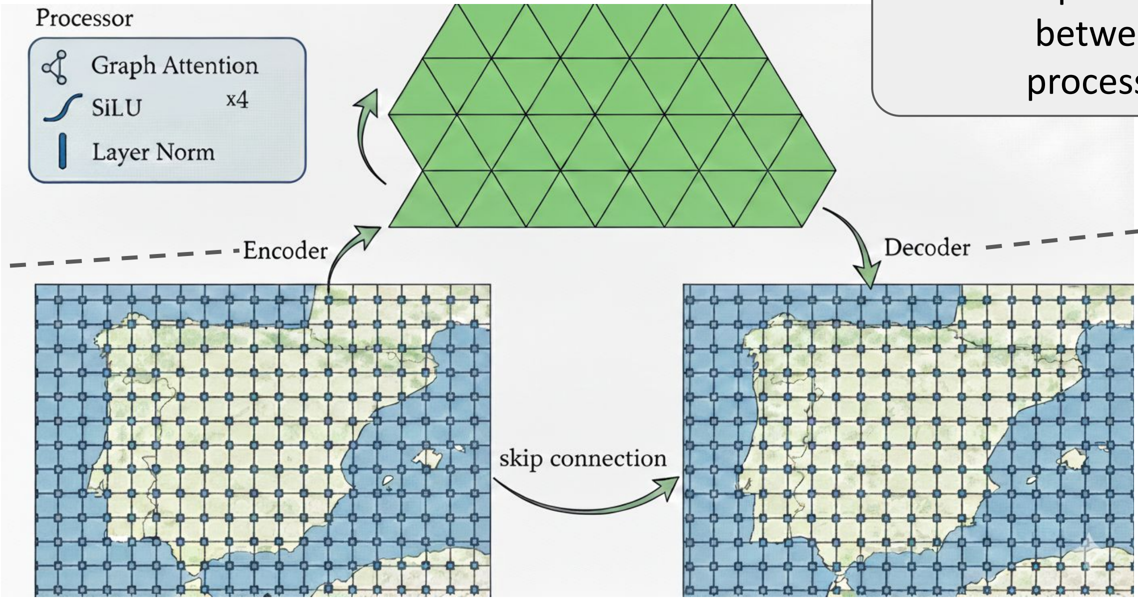

Encoder: MLP (1 hidden layer) + mean pooling

\+ skip connections between each processor layer

Decoder: MLP (1 hidden layer)

We predict **increments** $\Delta_t = \mathrm{conc}_{t+1} - \mathrm{conc}_t$, not state at t+1

Train with **L1Loss** with **cosine** learning rate scheduler for about 500 epochs with AdamW

GNN: the encoder

1

Build **regular** hexagonal mesh above the native MONARCH lon-lat curvilinear grid

2

Nodes are centers of hexagons

3

Dual graph makes **regular** **triangular** mesh

MLP
+
mean pooling

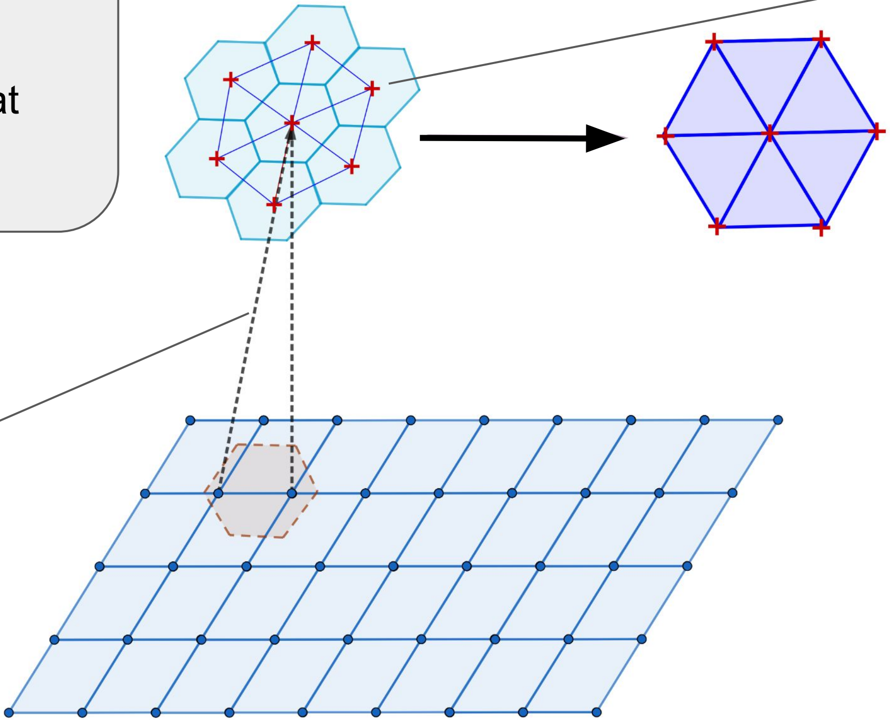

GNN: the processor

Use 4 layers of message-passing Graph Attention: $\mathbf{h}_i' = \sigma\left(\sum_{j \in \mathcal{N}(i) \cup i} \alpha_{i,j} \mathbf{W}_{right} \mathbf{h}_j\right)$

with the **attention weight (additive attention)**:

$$\alpha_{i,j} = \mathrm{softmax}_j\left(\mathbf{a}^\top \mathrm{LeakyReLU}\left(\mathbf{W}_{left}\mathbf{h}_i + \mathbf{W}_{right}\mathbf{h}_j\right)\right)$$

Learnable attention vector

Receiver state

Sender state

Similar to **self-attention mechanism** except only on **neighboring** nodes:
necessity to **stack** layers to propagate information deeper in the graph

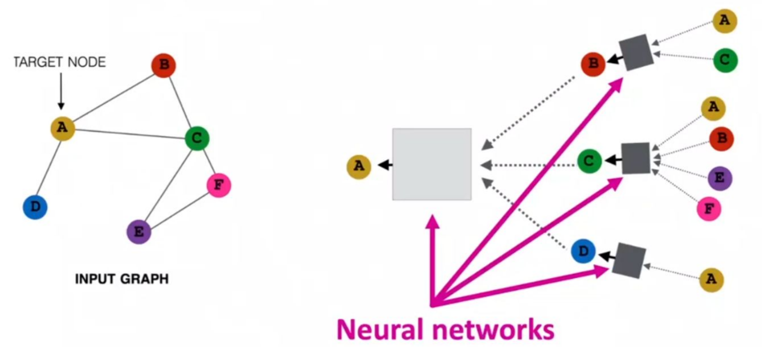

Results:
GNN versus UNet comparison

(focus on NO2 concentrations)

Spatial patterns

GNN

Mean	MB / nMB	MAE / nMAE	R^2	PCC
1.45e+00	-8.24e-03 / -1	1.72e-01 / 12	95.75	98.0

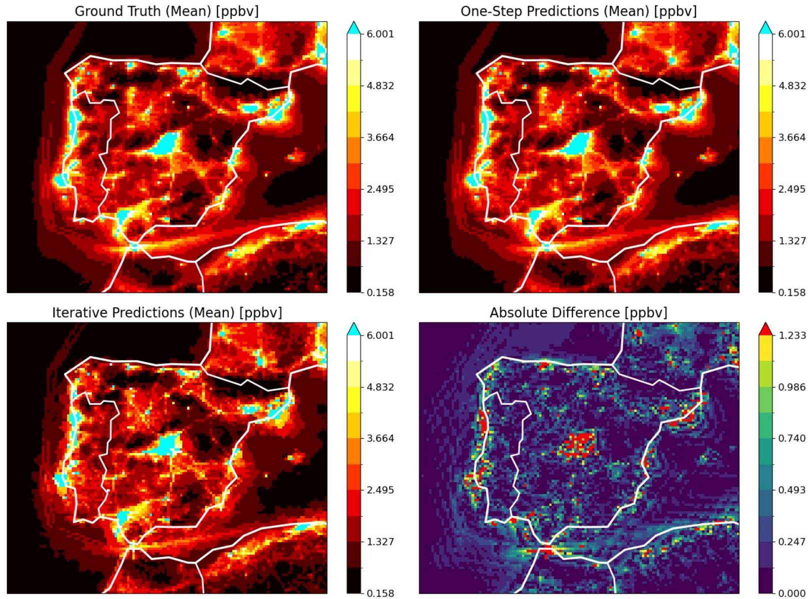

UNet

Mean	MB / nMB	MAE / nMAE	R^2	PCC
1.45e+00	5.62e-03 / 0	3.24e-01 / 22	97.62	98.7

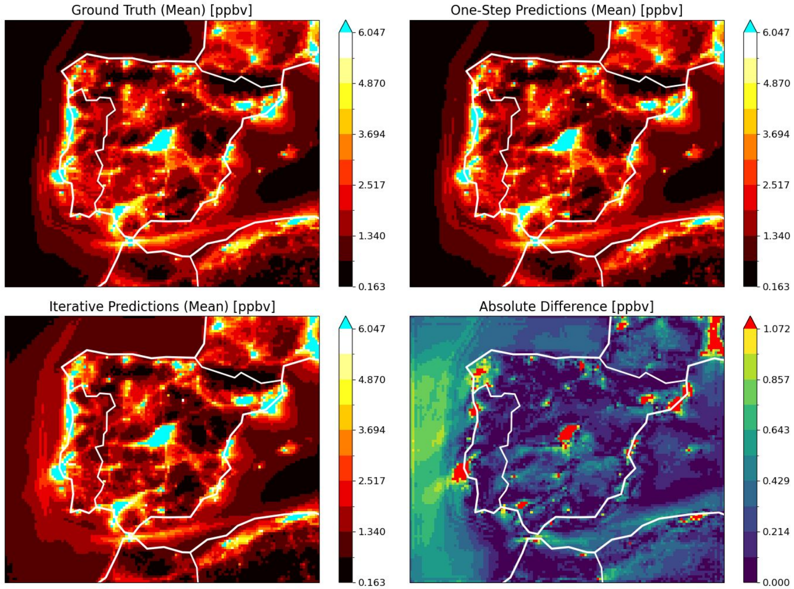

Performance of 1-hour predictions very good - **spatial patterns** partially learned well for autoregressive (2 weeks predictions)

Temporal patterns

15-days auto-regressive predictions initialised at the beginning of 4 seasons:

GNN

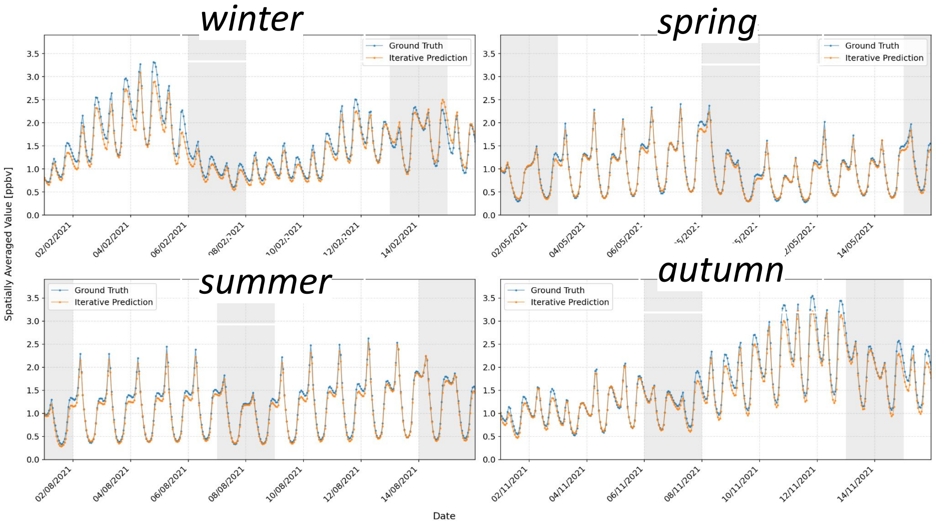

UNet

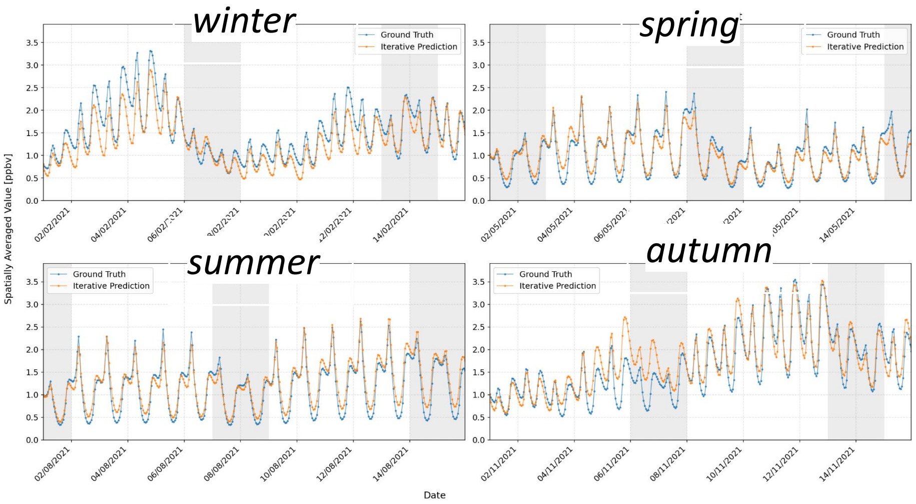

Reasonably good auto-regressive predictions with **GNN**, no critical divergence, variability well captured, less biased then **UNet**

Global statistics: Autoregressive rollout

Comprehensive scatterplots of auto-regressive predictions along 1-5, 6-10, and 11-15 days horizons:

GNN

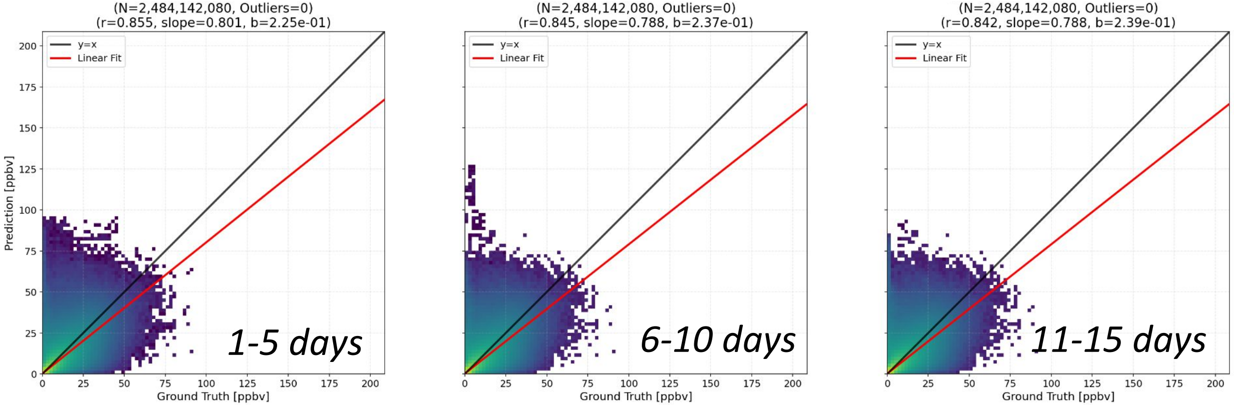

UNet

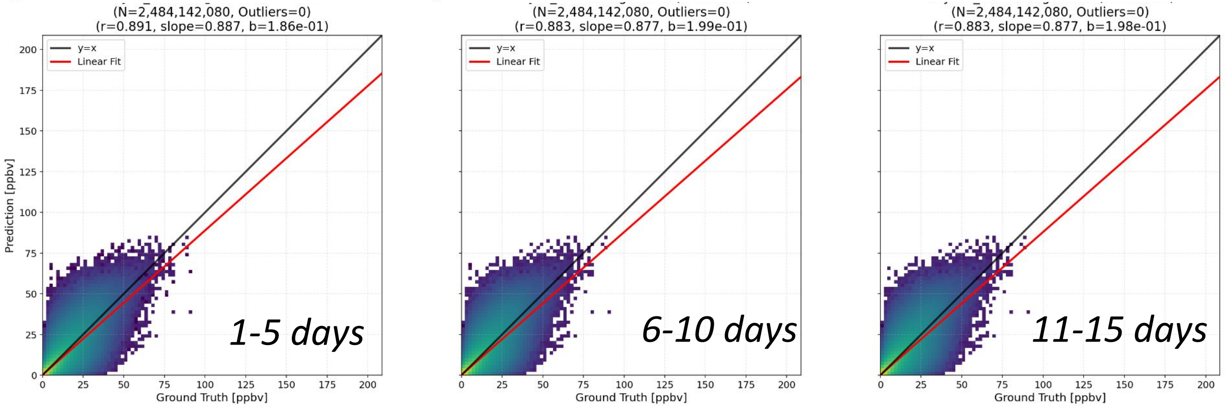

Overall **UNet** still shows slightly better auto-regressive performance, with correlation/slope around 0.88/0.88 (against 0.84/0.79 for **GNN**)

Recent trial: Swin Transformer - UNet

SwinTransformer

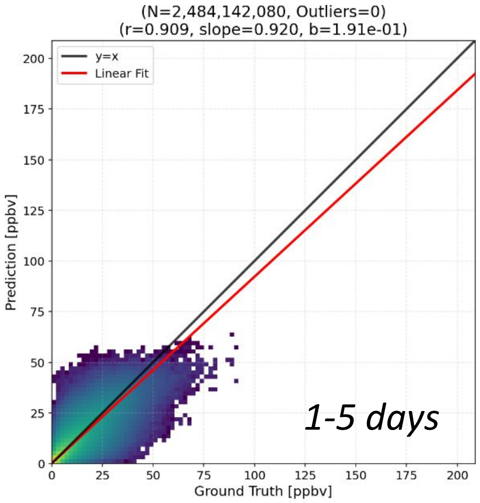

First results obtained with the recently tested **SwinTransformer** architectures show further improved performance compared to **UNet** (and **GNN**) with correlation/slope around 0.90/0.92

Conclusion and future prospects

Conclusion: promising first results on NO_2 — although less good on NO (not shown) — with reasonably good performance and stability of auto-regressive predictions

- More comprehensive tuning of these models
- More diverse training dataset size
- **Ongoing testing with more species** (e.g. O_3, PM10, PM2.5): more mitigated results at this stage, work in progress
- Ongoing implementation of **chemical boundary conditions** (important for longer-lived species and/or limited-size domains)
- Try architectural improvements for **GNN** (e.g. **Encoder** and **Decoder** as bipartite graph GNNs)

Acknowledgements

This presentation is part of the R&D project AIRE, funded by the grant "Proyectos de Investigacion en el ambito de la Inteligencia Artificial - 2025", Grant AIA2025-163866-C21 funded by MICIU/AEI/10.13039/501100011033. Additionally, the present work was funded through the RESPIRE project, which is part of the Recovery, Transformation and Resilience Plan (Plan de Recuperacion, Transformacion Resiliencia, PRTyR) funded by the European Commission to repair the harm caused by the Covid-19 crisis. We acknowledge the RES (AECT-2025-1-0012 and AECT-2025-2-0012) and EuroHPC (EHPC-AI-2024A06-068 and EHPC-REG-2025R01-219) for providing CPU and GPU computing resources. HP has received funding from the Ministerio de Ciencia e Innovacion (MCIN) and Agencia Estatal de Investigacion (AEI) through the Ramon y Cajal grant (RYC2021-034511-I, MCIN/AEI/10.13039/501100011033 and European Union NextGenerationEU/PRTR).

GOBIERNO DE ESPAÑA
MINISTERIO DE CIENCIA, INNOVACIÓN Y UNIVERSIDADES

GOBIERNO DE ESPAÑA
VICEPRESIDENCIA TERCERA DEL GOBIERNO
MINISTERIO PARA LA TRANSICIÓN ECOLÓGICA Y EL RETO DEMOGRÁFICO

MINISTERIO DE CIENCIA, INNOVACIÓN Y UNIVERSIDADES
Cofinanciado por la Unión Europea
AGENCIA ESTATAL DE INVESTIGACIÓN

Plan de Recuperación, Transformación y Resiliencia

Financiado por la Unión Europea
NextGenerationEU

Agencia Estatal de Meteorología

Thank you for you attention!