TOWARDS INVERTING SPANISH NOX EMISSIONS WITH TROPOMI OBSERVATIONS AND NEURAL NETWORKS

JAMES PETTICREW, HERVÉ PETETIN[1], MARC GUEVARA[1], ISIDRE MAS[1], ORIOL JORBA[1], CARLOS PÉREZ GARCÍA-PANDO[1-2]
[1] Barcelona Supercomputing Center (BSC), Barcelona, Spain
[2] ICREA, Catalan Institution for Research and Advanced Studies, Barcelona Spain
james.petticrew@bsc.es
herve.petetin@bsc.es

INTRODUCTION

Air pollution estimates represent key inputs in computer models for assessing air quality. They are also important in the evaluation of pollution control policies. In the last decade, neural networks have demonstrated exceptional ability to model complex spatiotemporal data, while advances in our ability to observe the earth's atmosphere using satellites have enabled the collection of high-resolution atmospheric composition data in near real-time. These developments open up opportunities to combine the predictive power of neural networks with satellite observations to deliver rapid and accurate estimates of pollutant emissions in near real-time. Chemical weather prediction models offer insights into the forward relationship between emissions and atmospheric composition, and some studies are already suggesting that neural networks might be able to estimate with reasonable predictive skills the chemical concentrations obtained from these physics-based models. While the forward mapping is well-defined, the inverse mapping—from atmospheric composition to emissions—is not. Our objective is ultimately to exploit neural networks to predict emissions from atmospheric composition. This presents challenges, as we will show in our presentation. Here, we will present preliminary results of our study on training neural networks to predict Spanish NOx (NO and NO2) emissions using data from Barcelona Supercomputing Centre's chemical weather prediction model NMMB-MONARCH and emission model HERMES. We used a UNet network structure to encode the inverse mapping, and trained the networks using the popular ADAM optimiser.

METHODOLOGY

DATA

Training data were generated by running HERMESv3 to produce hourly anthropogenic emission fields on a 121×141 rotated lat-lon grid over the Iberian Peninsula, based on the CAMS-REG-v4 inventory. Five emission scenarios were created by randomly perturbing sector-level magnitudes by factors of 0–2. These, combined with ERA5 meteorological fields, drove five MONARCH-NMMB CTM simulations spanning 2021–2024. Ocean, fire, biogenic, and natural soil NO emissions were prescribed from CAMS-GLOB-OCEAN-v11, GFASv12, and MEGAN v2.04, respectively.

From the CTM output, we extracted fields onto a reduced 96×96 regular lat-lon grid: (i) daily-mean near-surface meteorology (averaged over the two lowest pressure levels), (ii) daily total column emissions, and (iii) synthetic TROPOMI-like NO_2 total column densities at 13:00 UTC. The emission fields were the regression targets for the neural network and the meteorological fields and NO2 column densities form part of the predictors.

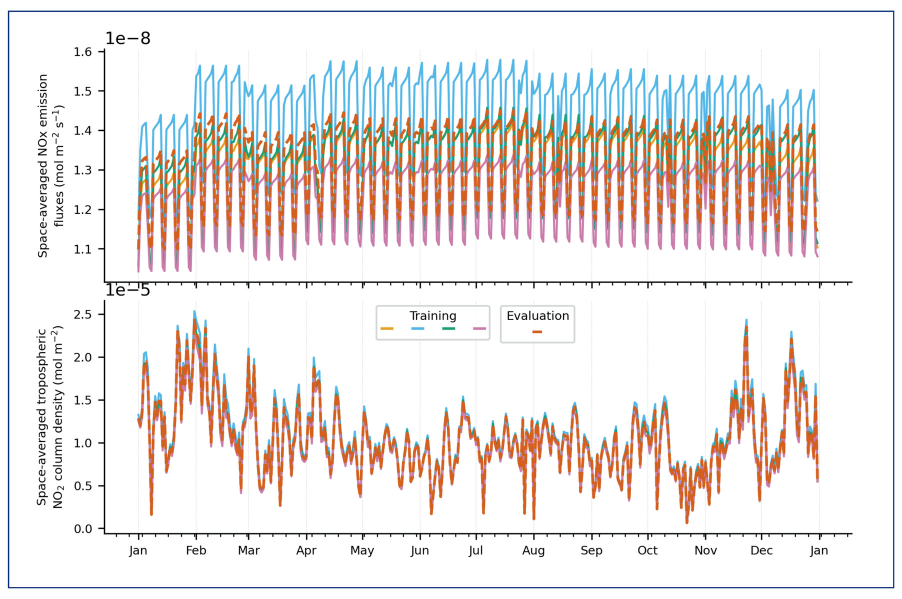

Figure (a): One year (2023) of NOx emissions and concentrations from our training data, averaged over space. The top panel shows the five perturbed emissions scenarios, and the bottom panel shows the corresponding atmospheric column densities of NO2 over the time period. We used one scenario (dashed line) for evaluation, and the remaining four scenarios for training.

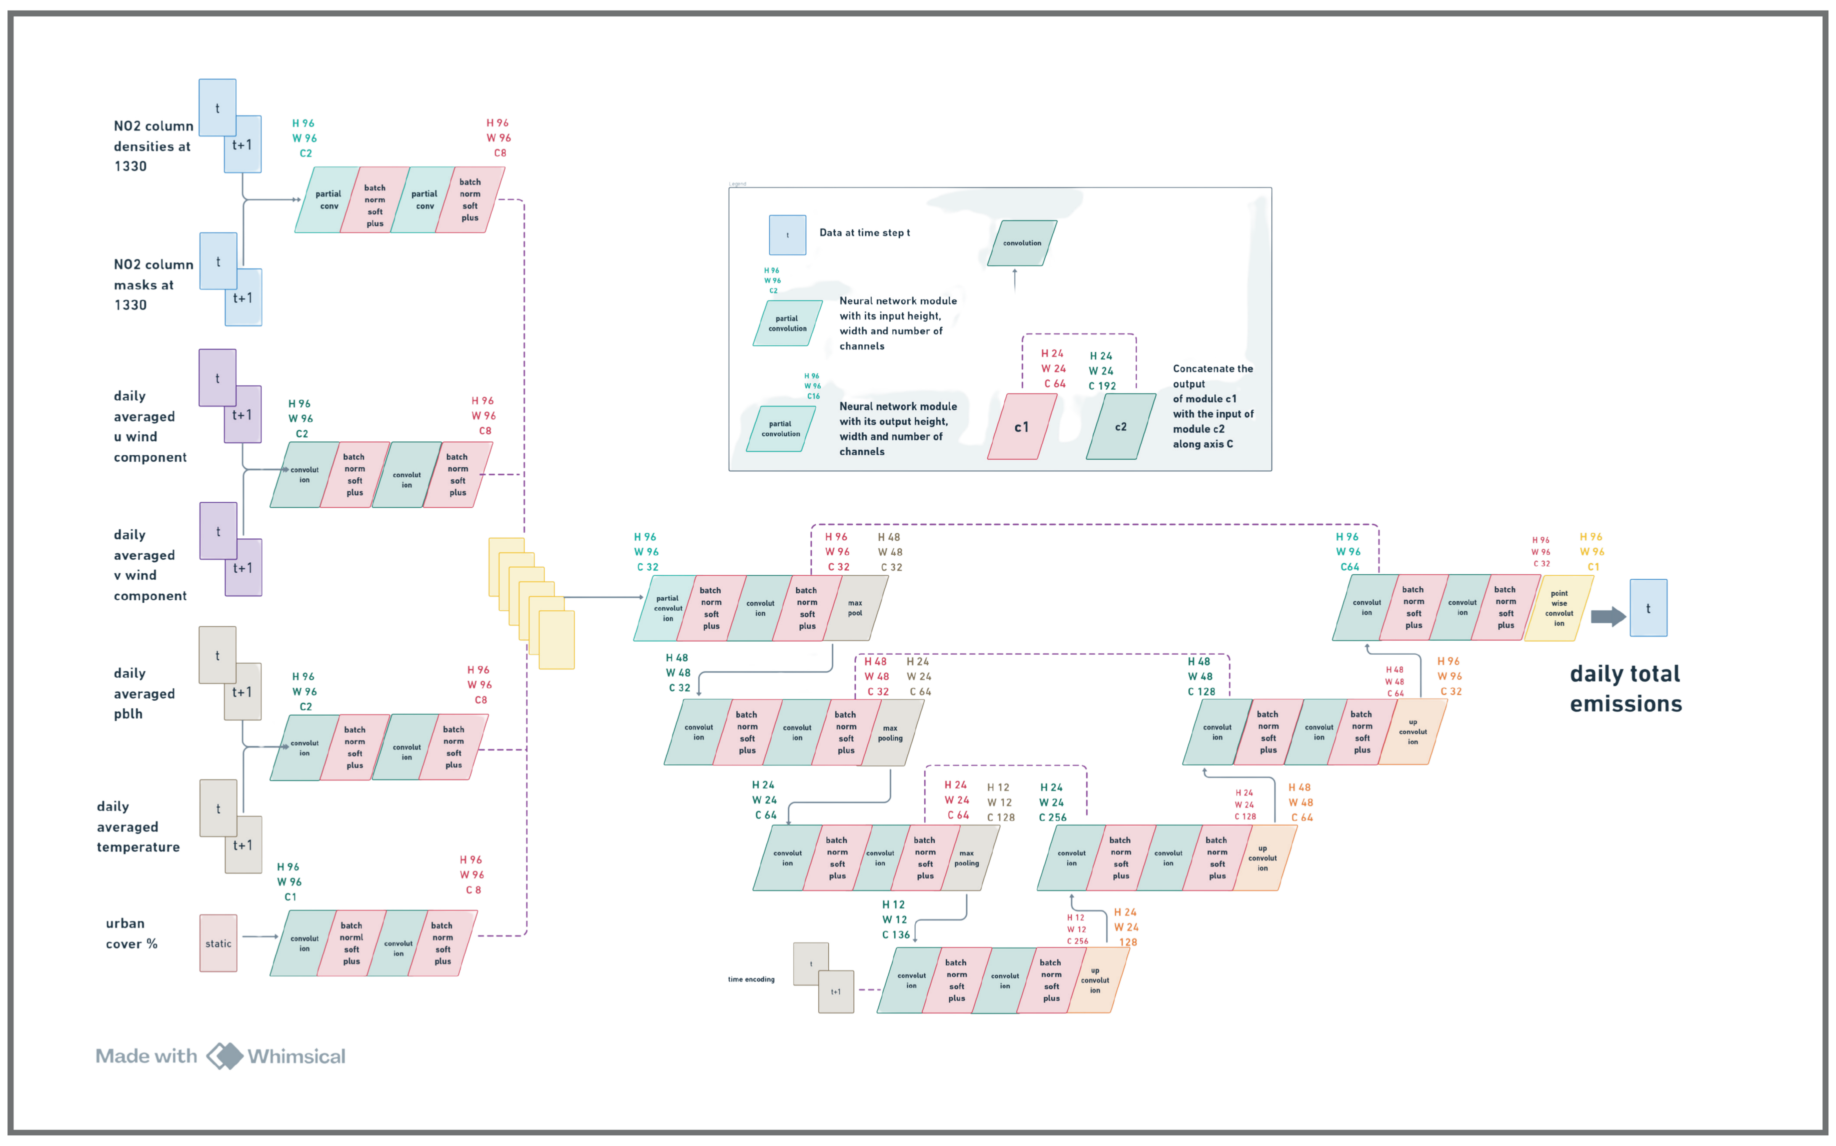

Figure (b): We train a U-Net to estimate daily NOx emissions across Spain from TROPOMI NO2 total columns at 13:30 UTC, daily-averaged meteorological fields, per-pixel urban cover fraction, and sinusoidal encodings of day-of-week and day-of-year. Each group of input variables is first processed by a dedicated stem of convolutional blocks before being concatenated and passed into the U-Net encoder. Because observed NO2_2 2 columns contain missing data due to cloud cover and retrieval failures, partial convolutions are used in the NO2_2 2 preprocessing stem to handle masked pixels. A single convolutional block at the U-Net entrance fuses the outputs of all stems, producing a complete feature map regardless of which NO2 pixels are valid. Sinusoidal time encodings are injected at the bottleneck, where the receptive field is global. Standard skip connections between encoder and decoder feature maps at each spatial scale allow the network to integrate both local detail and global context when predicting the emissions field.

TRANSLATING CTM DATA TO SATELLITE DATA

In order to train a model capable of predicting from TROPOMI-observed atmospheric NO2, we first interpolate the model NO2 mixing ratio profile onto the TROPOMI pressure grid, then apply the TROPOMI averaging kernel to the interpolated CTM profile, using the TROPOMI a priori profile and pressure weighting function. We then integrate over the column using the pressure weighting function to get a total column density, before applying a qa filter to retain only pixels with valid retrievals.

TRAINING

We trained our model on 4 Platinum 8460Y+ cores, each with 4xH100 Nvidia GPUs. We trained for 2000 epochs, with a ReduceLROnPlateau learning rate scheduler with patience of 100 epochs and a reduce factor of 0.5. We tuned our learning rate on a log grid to find an optimal value of 1e-3. We used the ADAM optimiser and a mean-squared error loss. We train the model on four of the perturbed emission scenarios, and tune on the remaining +scenario. On our evaluation set we obtained a normalised mean bias of -2.40%, a normalised MAE of 11.00% and R2 of 97.1%

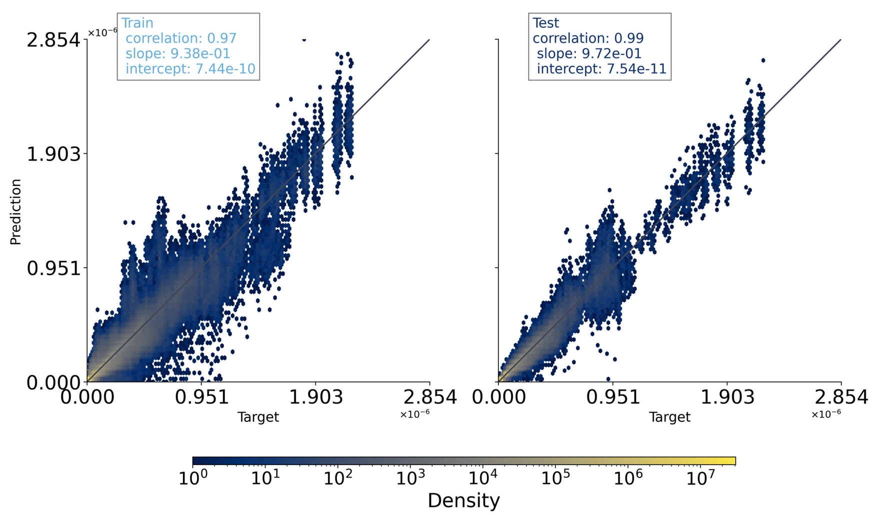

Figure (c): Correlation scatter plots between target and predicted emissions from our trained network.

RESULTS AND DISCUSSION

(d) Valladolid

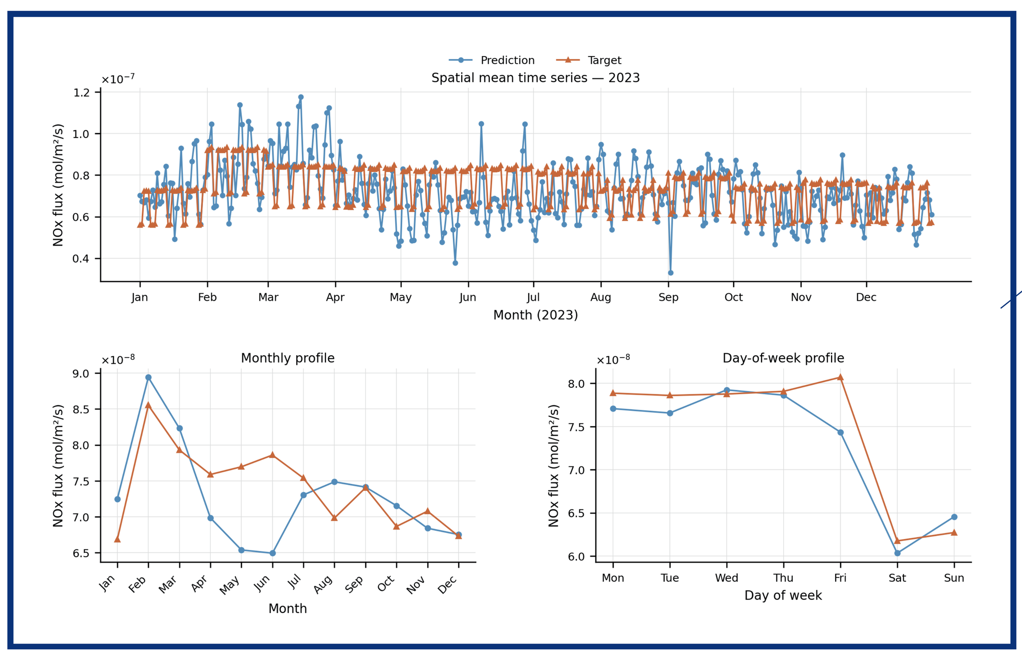

Figure (d): Predicted and target NOx emissions in Valladolid throughout 2023 - time series (top pane), monthly profile (bottom left) and weekly profile (bottom right) for the evaluation scenario. In the cities, NOx emissions are dominated by the "weekend effect", which the model captures well on average, but underpredicts in the summer on average.

(e) Madrid

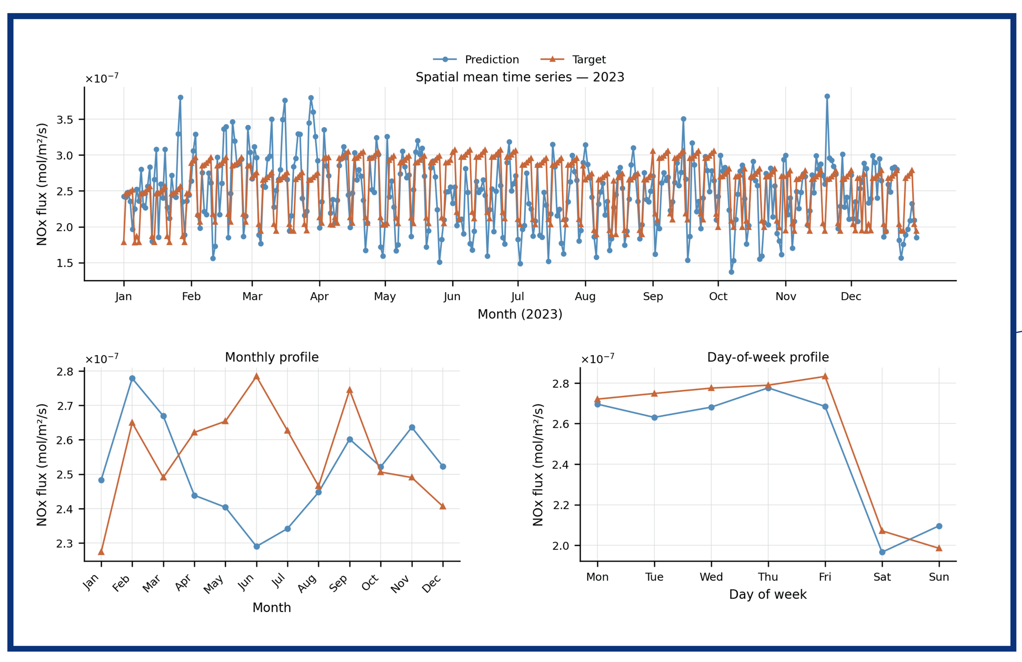

Figure (e): Predicted and target NOx emissions in Madrid throughout 2023 - time series (top pane), monthly profile (bottom left) and weekly profile (bottom right) for the evaluation scenario. In Madrid NOx emissions are around 3x higher than in Valladolid. In Madrid, emissions are expected to rise in the summer, but the model predicts that they fall.

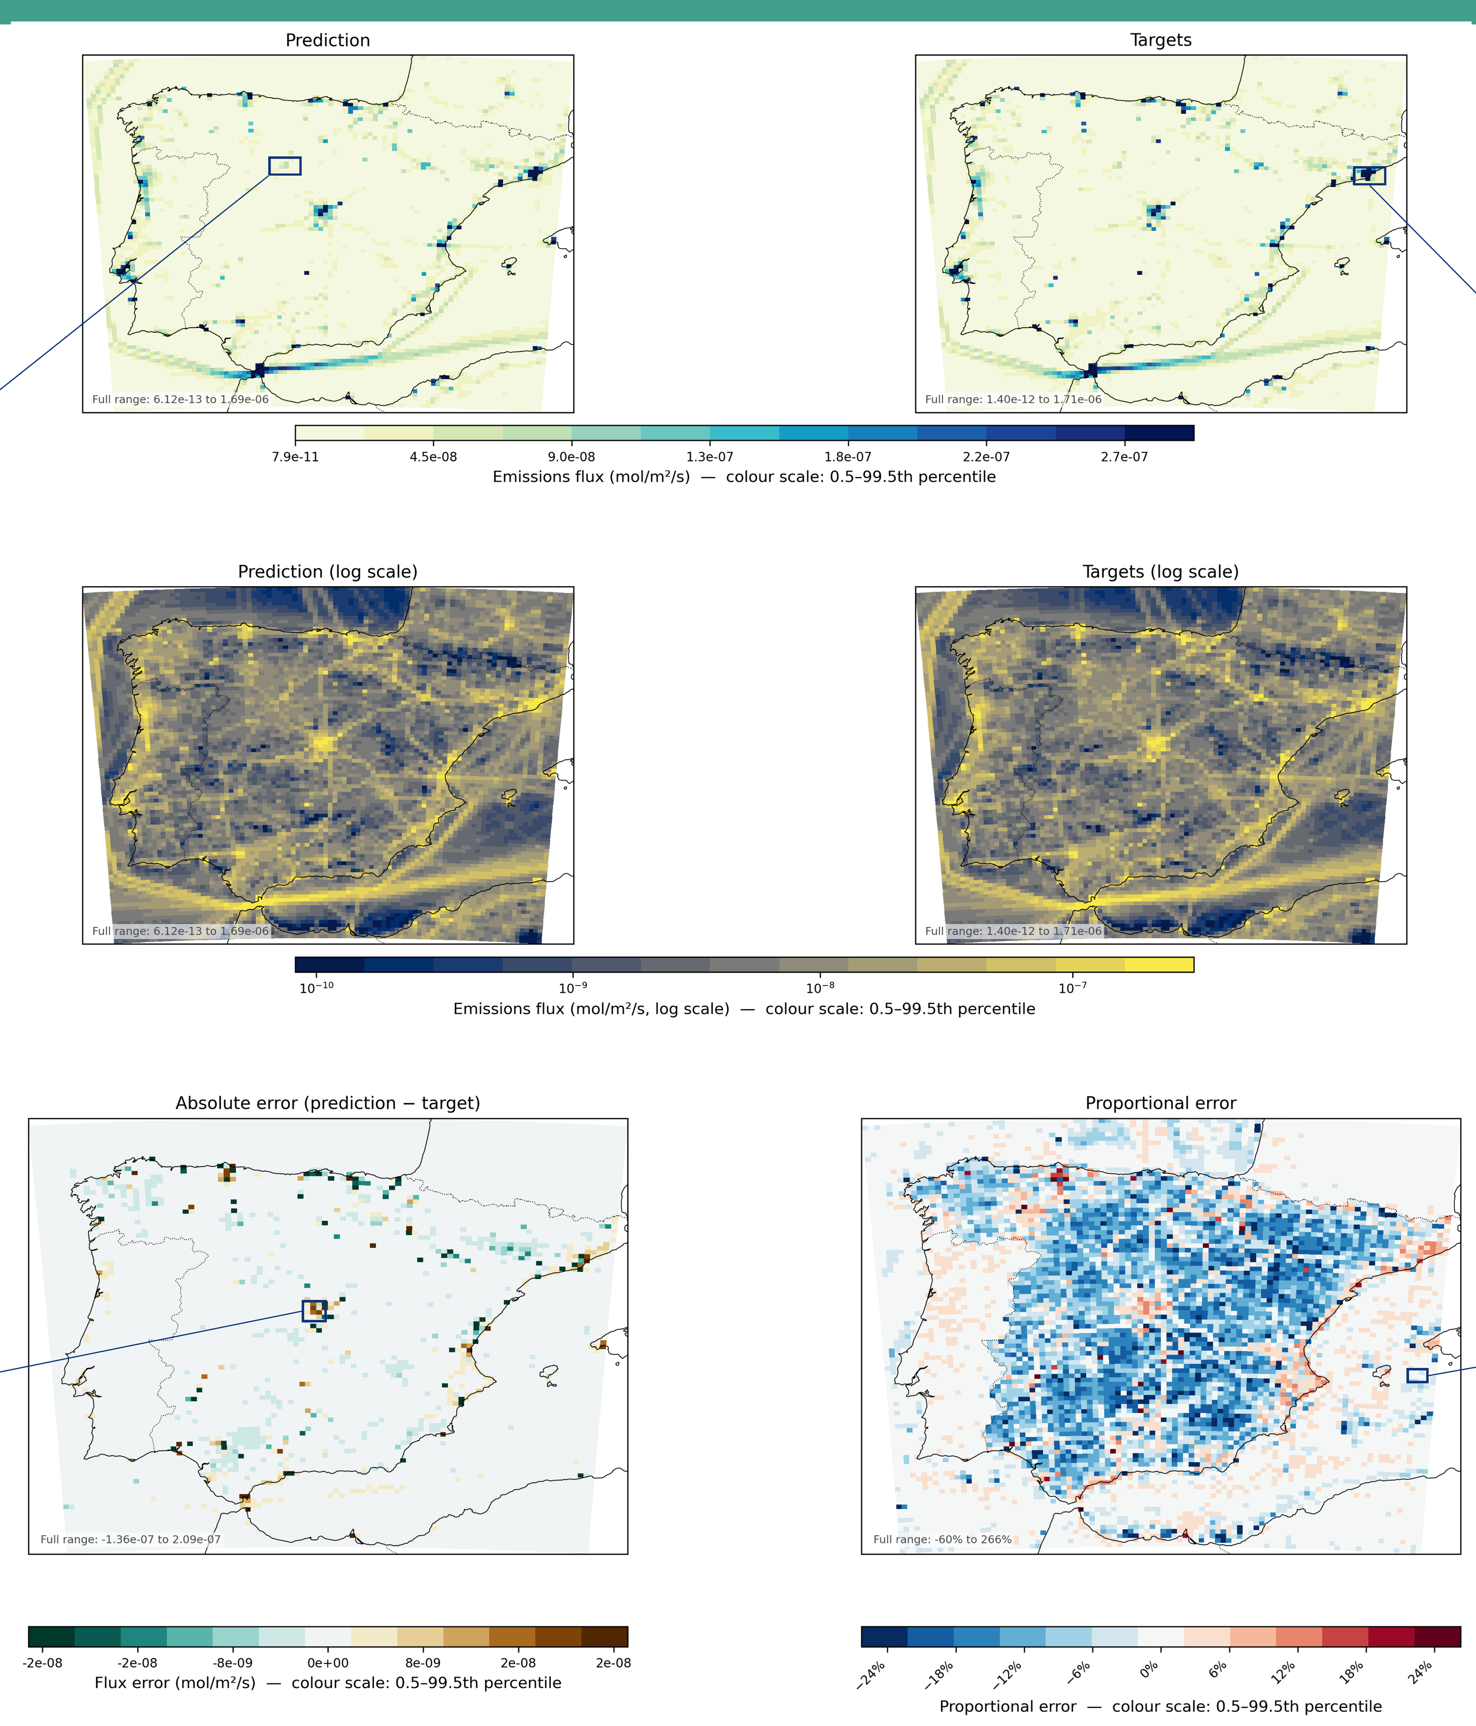

(f) Time-averaged emission fluxes over the Iberian Peninsula - evaluation set

Figure (f): These maps are averaged over the 2021-2024 training period. The top row shows averaged predicted and target NOx fluxes on a linear scale, the second row shows them on a log scale, and the third row shows absolute and relative differences between the ML model predictions and the evaluation set targets.

(g) Barcelona

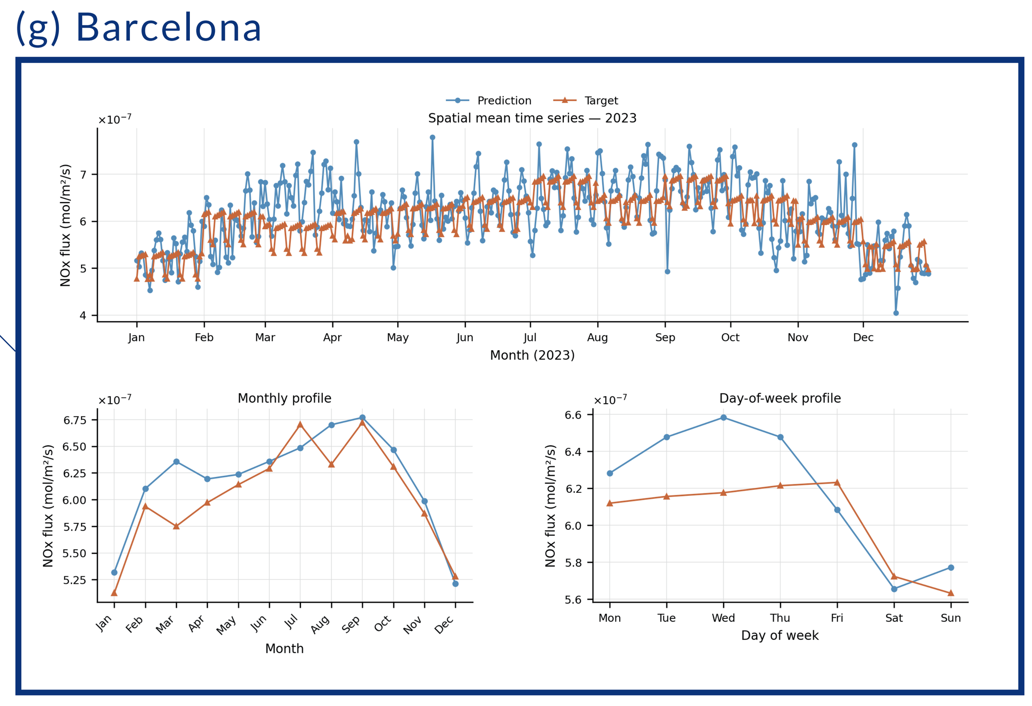

Figure (g): Predicted and target NOx emissions in Barcelona throughout 2023 - time series (top pane), monthly profile (bottom left) and weekly profile (bottom right) for the evaluation scenario. In contrast to in Valladolid and Madrid, the model well captures the monthly profile but overpredicts the weekly profile.

(h) Balearic Sea (near Ibiza)

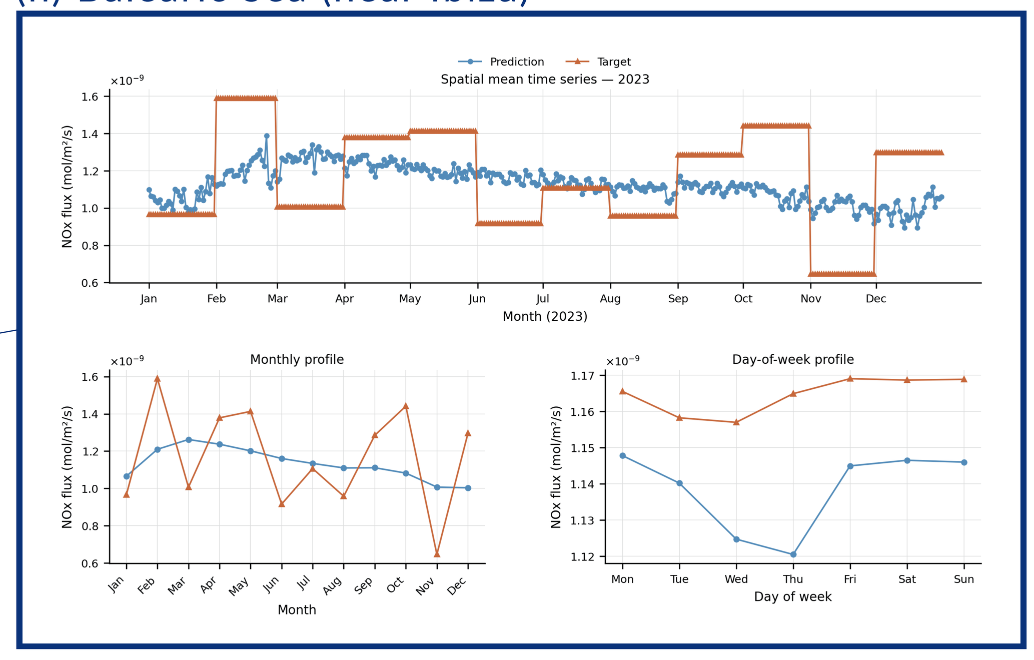

Figure (h): Predicted and target NOx emissions in Balearic Sea throughout 2023 - time series (top pane), monthly profile (bottom left) and weekly profile (bottom right) for the evaluation scenario. Here, emissions are typically two orders of magnitude lower than in the cities, and the training data from the simulator only shows significant relative variability on the monthly scale, not captured by the model.

WHAT'S NEXT?

The ultimate aim of this project is to leverage this model to generate near-real time emission estimates from satellite observations, and to generate emission estimates to help in the validation of other, more detailed, emission inventory emissions techniques. Our next step will be to generate total NOx emissions estimates over the Iberian Peninsula using TROPOMI observations, and to analyse the implications of the results on existing emission inventories, given the performance of our model. We want to look at more powerful model architectures, which may imply generating more training data. We are interested in applying a similar methodology to hourly observations from Sentinel-4.

ACKNOWLEDGEMENTS

The project that gave rise to these results received the support of a fellowship from the "la Caixa" Foundation (ID 100010434). The fellowship code is LCF/BQ/D123/11990077. HP has received funding from the Ramon y Cajal grant (RYC2021-034511-I, MCIN / AEI / 10.13039/501100011033 and European Union NextGenerationEU/PRTR). We acknowledge the RESPIRE project, which is part of the Recovery, Transformation and Resilience Plan (Plan de Recuperación, Transformación y Resiliencia, PRTyR) funded by the European Commission to repair the harm caused by the Covid-19 crisis.

"la Caixa" Foundation
Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA — VICEPRESIDENCIA TERCERA DEL GOBIERNO — MINISTERIO PARA LA TRANSICIÓN ECOLÓGICA Y EL RETO DEMOGRÁFICO
RESPIRE AIR
Barcelona Supercomputing Center — Centro Nacional de Supercomputación