Enhancing air quality simulations with neural network-based resolution downscaling

Anthony Rey-Pommier[a], Enrico Pisoni[a], Philippe Thunis[a], Stefano Zauli-Sajani[a], Alexander de Meij[b], Maxime Beauchamp[c], Bertrand Bessagnet[d]

Background and aim

Air pollution is the second leading risk factor for mortality, primarily due to exposure to fine particles and nitrogen oxides (NO_x = NO + NO_2). High-resolution data for these pollutants are therefore crucial for exposure assessment and policymaking. However, the use of chemistry-transport models (CTMs), that evaluate pollutant concentrations based on emission inventories and weather data, is computationally intensive at high-resolution. To improve the resolution of such estimates, machine learning methods, especially deep learning, can be used to downscale CTM outputs. We present a super-resolution model based on the Residual Channel Attention Network[1] (RCAN), trained on outputs of the CTM EMEP[2]. We evaluate the model for $PM_{2.5}$ (fine particles with diameter below 2.5 µm) and NO_2 concentrations over Europe.

Input and target variables, domain and validation metrics

Input variables include hourly emissions of primary particulate matter, NO_x, SO_x, NH_3 and NMVOCs taken from CAMS_v6.1 at low resolution (LR, ~10×5 km^2) and high resolution[3] (HR, ~2×1 km^2), horizontal wind from the IFS (LR), as well as concentrations (LR) and the vertical turbulent diffusion coefficient (LR and HR) from EMEP outputs. The model downscales concentration fields over a large domain comprising most of Europe for January 2019. To ensure efficient learning and generalization, we apply a patch-based training strategy, dynamically sampling spatial patches at each epoch. During inference, overlapping patches are processed and merged using central cropping to minimise edge artefacts. The model is trained over January 1–25, and validated over January 26–31, using two scenarios: a Base Case representing the state of the atmosphere with emissions of 2015, and a second scenario where a 50% reduction in all pollutant emissions is applied (-50%ALL). Performance is evaluated using standard metrics such as the normalised Root Mean Square Error (RMSE) and the coefficient of determination (R^2).

Model structure

We use the Residual Channel Attention Network (RCAN), a deep learning architecture designed for super-resolution tasks. Its Residual-in-Residual (RIR) blocks and Channel Attention (CA) mechanisms allow the model to focus on informative features and reconstruct fine-scale spatial patterns more effectively than traditional architectures. Its depth and attention capabilities make this architecture well-suited for downscaling pollutant fields over large and heterogeneous domains. We avoid generative models (e.g., GANs, diffusion models) due to their stochastic nature and limited suitability for deterministic, physically constrained tasks like pollutant downscaling. The objective is not to generate plausible patterns but to reconstruct known high-resolution structures from coarse-resolution CTM outputs.

Explanatory variables:
Surface $PM_{2.5}$ or NO_2 concentration (LR)
10m horizontal wind (LR)
Surface O_3 concentration (LR)
Surface NH_3 concentration (LR)
Surface NMVOC concentration (LR)
Surface SO_2 concentration (LR)
Surface NO_2 concentration (LR)
NH_3 emissions (HR & LR)
NO_x emissions (HR & LR)
Primary PM emissions (HR & LR)
SO_x emissions (HR & LR)
VOC emissions (HR & LR)
Vertical diffusion coefficient (HR & LR)

Target variable:
Surface $PM_{2.5}$ or NO_2 concentration (HR)

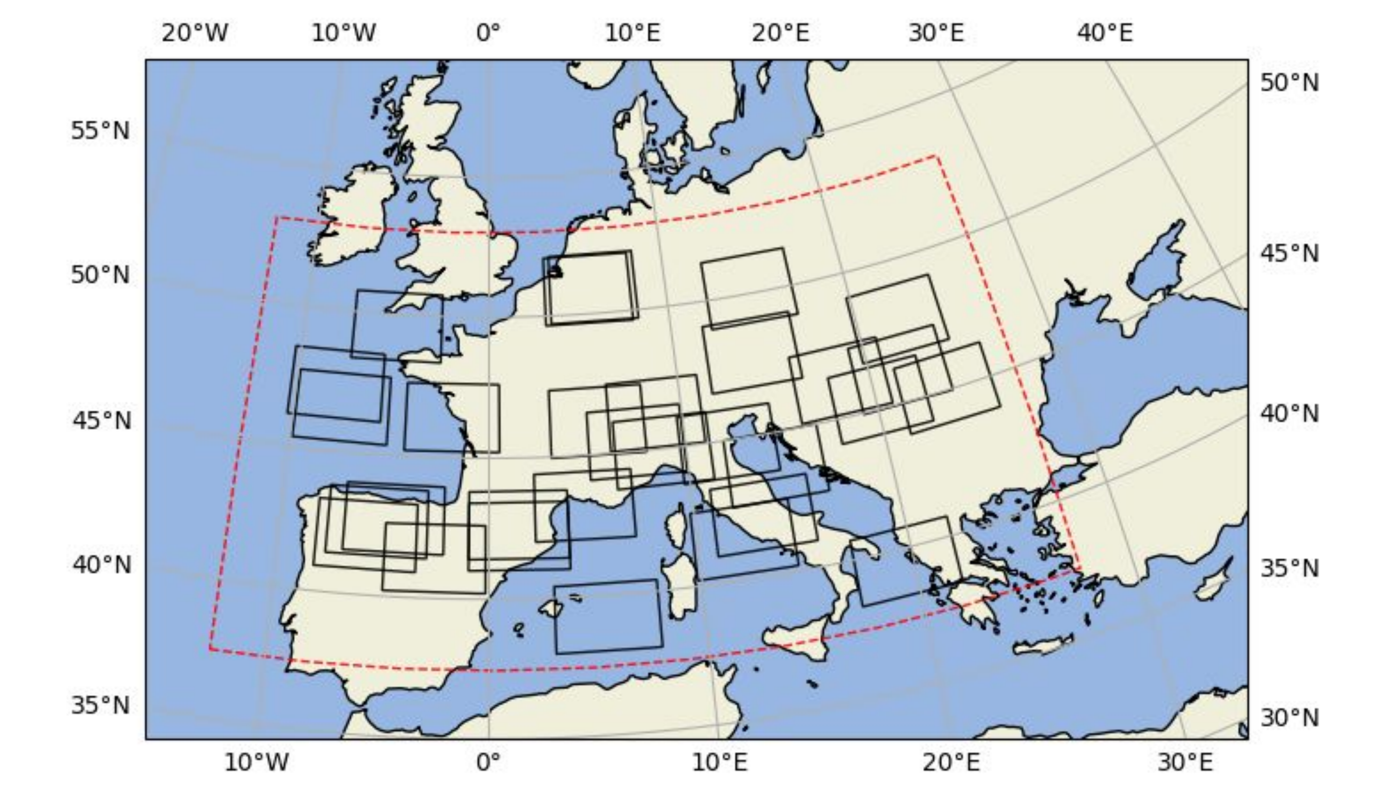

Used European domain (dotted line) and patch strategy for the training process.

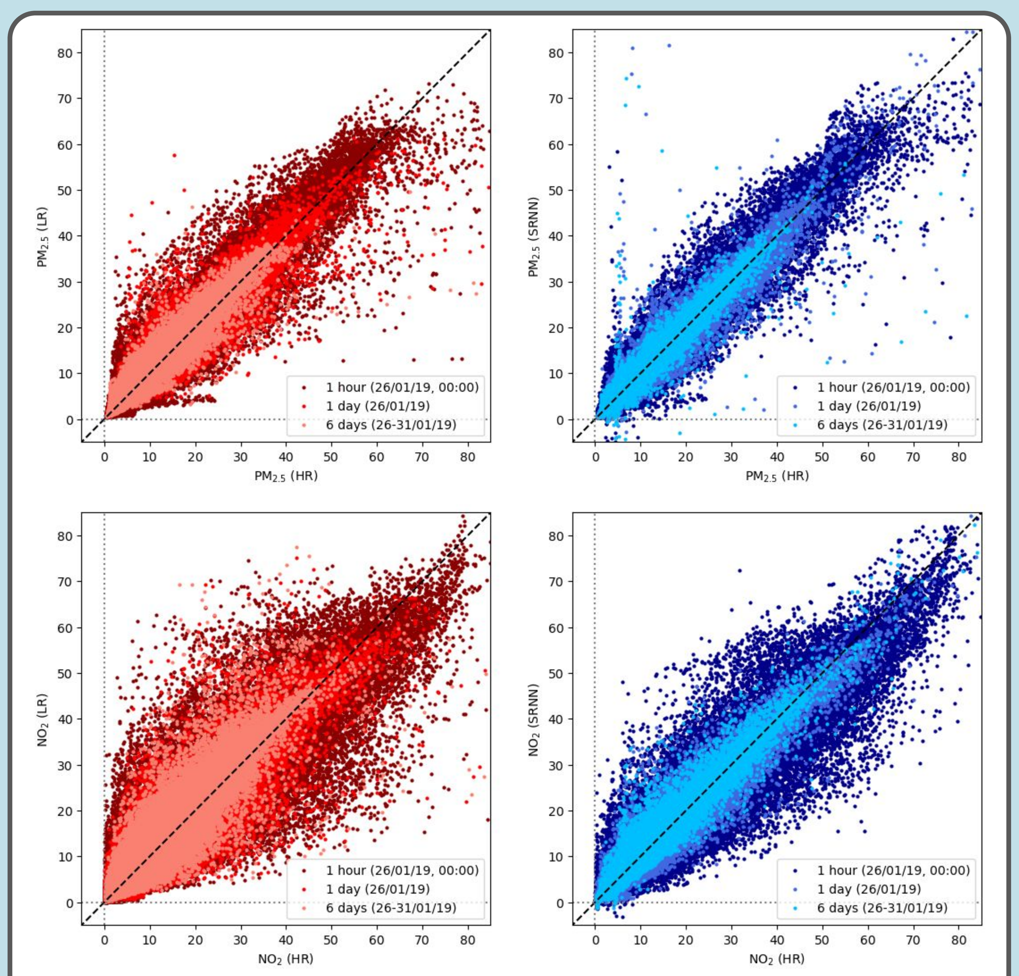

Comparison of low-resolution concentrations (LR, left) and super-resolution (SR, right) outputs with high-resolution (HR) concentrations for $PM_{2.5}$ (top) and NO_2 (bottom). Each pixel in the domain is represented by three points corresponding to three temporal averagings (1 hour, 1 day, 5 days).

Averaging horizon

Increasing the temporal averaging horizon significantly improves agreement between high-resolution simulated concentrations and super-resolution outputs, particularly within the first 24 hours. Between both quantities, the increase in performance metrics is sharp up to around the 1-day averaging mark, beyond which it becomes gradual. A daily resolution thus appears to be the optimal trade-off between preserving temporal detail and achieving acceptable performance.

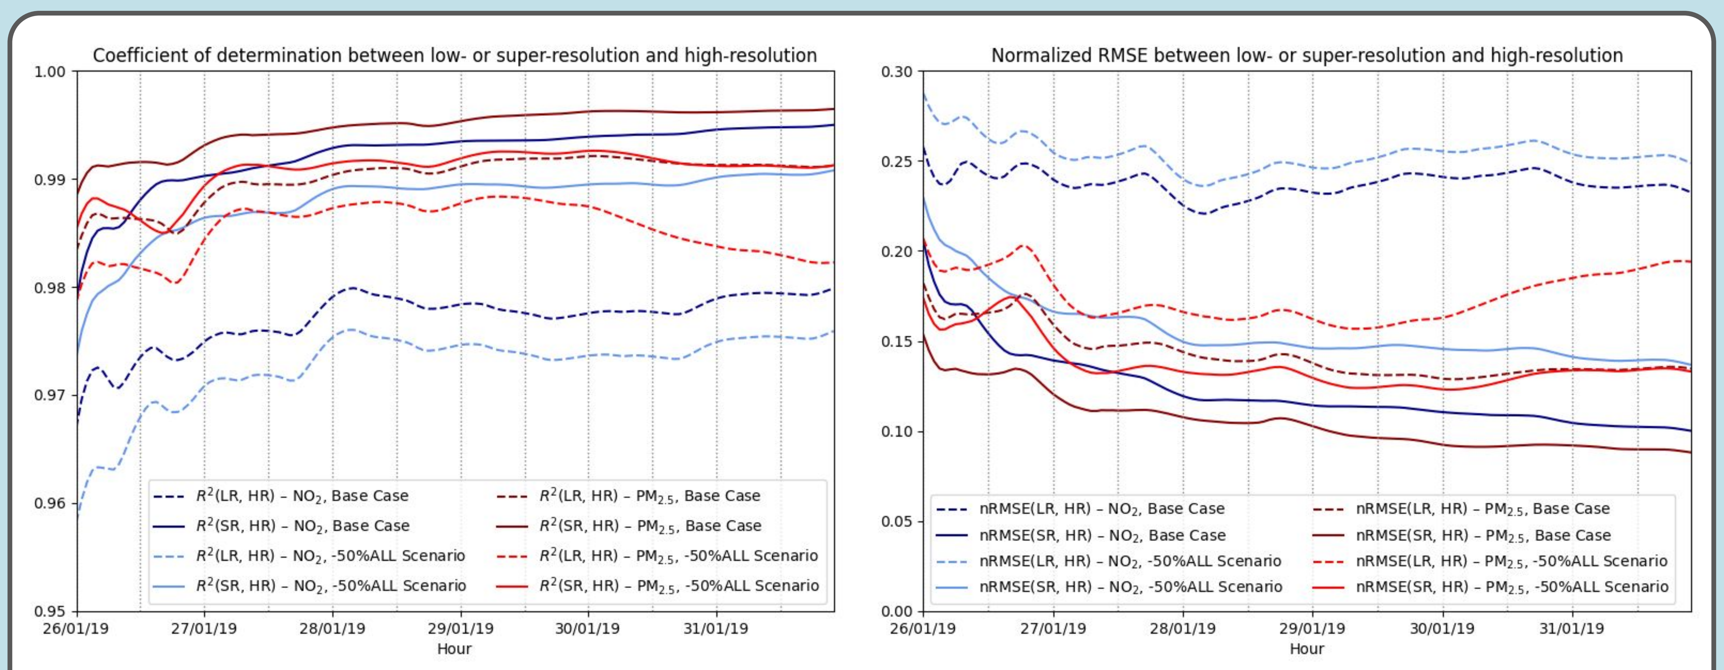

Evolution of the coefficient of determination (left) and normalised RMSE (right) between low-resolution concentration fields or super-resolution outputs and high-resolution simulations for NO_2 and $PM_{2.5}$, as a function of increasing temporal averaging horizon. Results are shown for the Base Case and the 50%ALL emission scenario, where all emissions are halved.

Assessment of results

The super-resolution model is able to reproduce fine-scale details from low-resolution inputs that usually appear in high-resolution simulations. This is particularly the case in large cities and industrial areas, where emission budgets are a significant driver of pollution patterns. High values of performance metrics demonstrate a high level of accuracy for both the Base Case and the -50%ALL scenario, where passive and active species are reduced. This suggests that the model successfully replicates different chemical regimes, transforming primary emissions into $PM_{2.5}$ and NO_2. This also indicates that the spatial and temporal variability in input data, especially emissions, is sufficient for training a reliable super-resolution model. This model could be used, after minor improvements and treatment of outliers (e.g. the presence of negative concentrations and isolated peaks), as an alternative to CTM simulations at daily temporal resolution and high spatial resolution. Results could also be improved by studying other pollutants, such as ozone, and by training the network directly on scenario differences as the target variable, rather than absolute concentrations. However, this approach would come at the cost of reduced flexibility, as it would require the availability of both a Base Case and an emission reduction scenario. This highlights the potential of the RCAN network to develop valuable tools for air quality research, enabling the exploration of various emission reduction strategies with increased precision and efficiency.

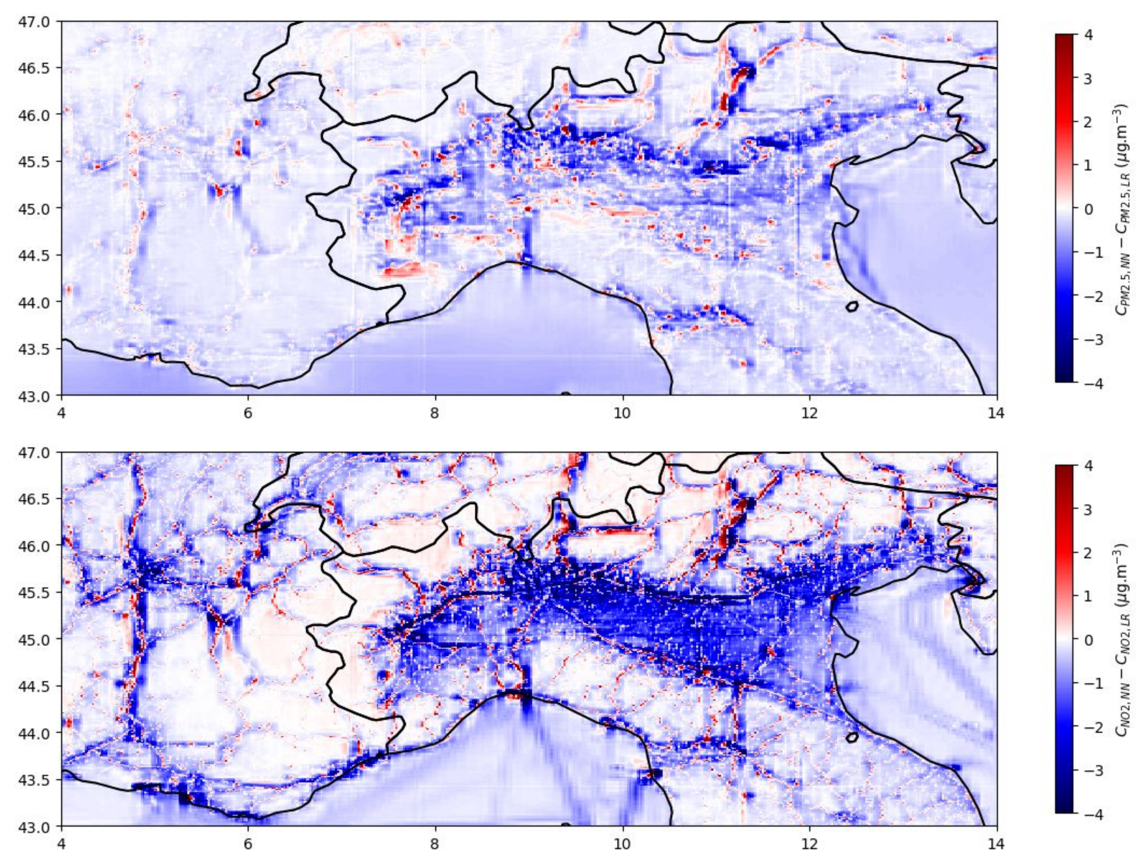

Concentration enhancement (difference between super-resolution outputs and low-resolution data), averaged for 26-31/01/19, for $PM_{2.5}$ (top) and NO_2 (bottom), for the Alps, the Po and the Rhône Valley.

Conclusion

This study presents a machine learning-based approach for downscaling $PM_{2.5}$ and NO_2 concentrations over Europe. The method achieves correct generalization across diverse regions and conditions. A key strength is its ability to reproduce fine-scale spatial structures without retraining for specific regions or scenarios. The results it provides are promising for delivering rapid air quality mapping and scenario screening[4], in particular for reflecting local emission reduction scenarios, which are crucial in air quality policymaking. Future work could explore extending the approach to other pollutants, such as ozone, and explore alternative models that use scenario differences rather than absolute concentrations as target variables.

References :

- [1] Zhang et al., 2018: Image super-resolution using very deep residual channel attention networks. Proceedings of the European conference on computer vision (ECCV), 286-301.
- [2] Simpson et al., 2012: The EMEP MSC-W chemical transport model–Technical description. Atmospheric Chemistry and Physics, 12(16):7825–7865.
- [3] Bessagnet et al., 2023: A simple and fast method to downscale chemistry transport model output fields from the regional to the urban/district scale. Environmental Modelling & Software, 164:105692
- [4] Thunis et al., 2016: On the design and assessment of regional air quality plans: The SHERPA approach. Journal of Environmental Management, 183:952–958.

EU Science Hub
joint-research-centre.ec.europa.eu

Joint Research Centre

[a] European Commission, Joint Research Centre, Directorate Energy, Mobility and Climate, Clean Air and Climate Unit, Ispra, Italy
[b] MetClim, Varese, Italy
[c] Danish Meteorological Institute, Copenhagen, Denmark
[d] LMD/IPSL, École Polytechnique, Institut Polytechnique de Paris, ENS, PSL Research University, Sorbonne Université, CNRS, Palaiseau, France

Contact: Anthony Rey-Pommier, anthony.rey-pommier@ec.europa.eu