



HAL
open science

State estimation of urban air pollution with statistical, physical, and super-learning graph models

Matthieu Dolbeault, Olga Mula, Agustín Somacal

► **To cite this version:**

Matthieu Dolbeault, Olga Mula, Agustín Somacal. State estimation of urban air pollution with statistical, physical, and super-learning graph models. 2024. hal-04455853

HAL Id: hal-04455853

<https://hal.science/hal-04455853>

Preprint submitted on 13 Feb 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

State estimation of urban air pollution with statistical, physical, and super-learning graph models

Matthieu Dolbeault, Olga Mula, and Agustín Somacal

Abstract

We consider the problem of real-time reconstruction of urban air pollution maps. The task is challenging due to the heterogeneous sources of available data, the scarcity of direct measurements, the presence of noise, and the large surfaces that need to be considered. In this work, we introduce different reconstruction methods based on posing the problem on city graphs. Our strategies can be classified as fully data-driven, physics-driven, or hybrid, and we combine them with super-learning models. The performance of the methods is tested in the case of the inner city of Paris, France.

1 Introduction

1.1 Background and motivation

Data-driven estimations are becoming increasingly relevant and widespread as the volume and heterogeneity of available data increases. A fundamental challenge is to build numerical methods for which one can estimate how optimally they exploit the given information. The present paper addresses some essential computational aspects connected to this question. More specifically, our goal is to reconstruct a state u of a physical process, for which we have at hand very heterogeneous sources of data coming from direct partial observations of u , from quantities related to u , and from the knowledge that the physics can be modelled by a Partial Differential Equations (PDE).

Assume that u belongs to some Banach space U of potentially infinite dimension, with associated norm $\|\cdot\|_U$, and that all the available information is given by an element x_u from some abstract metric space \mathcal{X} . Our goal is thus to build a mapping $A : \mathcal{X} \rightarrow U$ such that $A(x_u)$ approximates u at best, in the sense that the approximation error

$$e(A, u) = \|u - A(x_u)\|_U \tag{1.1}$$

is as small as possible, for any configuration (u, x_u) of the system. In practice, finding the optimal map A is not feasible, and various suboptimal reconstruction techniques have been proposed, each of them having its own virtues and drawbacks: statistical approaches such as BLUE [1] and kriging [2], model order reduction of parametric PDEs [3, 4], or more recently approximations by neural networks and machine learning strategies [5, 6]. Since all of these strategies are sub-optimal, and each one is based on different a priori assumptions, one should not make the methods compete against each other, but rather collaborate with each other to enhance their respective strengths. This leads naturally to explore approaches based on ensemble super-learning [7, 8, 9] as we consider in the present work.

1.2 Urban air pollution modeling

There are numerous applications in which one is confronted with the above state estimation problem. As a guiding example, we consider in this paper the real-time reconstruction of urban

pollution fields. Beyond the relevance of such a task to limit environmental and health risks in the city, pollution state estimation is an excellent example where collaborative, super-learning methods are required. This is because the problem accumulates several difficulties that make the reconstruction challenging for most common reconstruction methods. Among the issues, we may mention the following:

- Scarcity of pollution measurements: The amount of reliable sensor devices measuring pollutant concentrations is often limited, and the measurements are usually taken at fixed locations. As a result, reconstruction methods based solely on these measurements lack spatial resolution, and exhibit huge uncertainties in regions without sensors.
- Heterogeneous data: In addition to the pollutant measurements, other sources of relevant information are available such as traffic estimations in each street, wind speed, topography, temperature, etc. However, it is not obvious how to meaningfully combine this data to enhance the estimation. Some attempts have been tried in [10] through the use of Gradient Boosting Machines and Universal Krigging, with positive results for the estimation of PM10 particles in the city of Barcelona.
- Lack of training data: Even when incorporating other sources of information, the available data may be insufficient, noisy or hardly correlated to the pollution levels we wish to characterize. Purely data-driven models greatly suffer from these impediments in their training phase.
- Complexity of the physical problem: The equations governing the dispersion of pollutants in the atmosphere are nonlinear, with turbulent effects at the street scale, thus imposing a fine spatial resolution, at least near the sensor stations [10]. On the other hand, the computational domain is of the size of a city, making it prohibitively expensive to solve a full model like 3D Navier-Stokes equations.
- Parameter calibration: Reduced models use effective parameters, which account for large-scale averages of local effects, in order to alleviate the requirements on the resolution. However these parameters must be calibrated based on the available data or preliminary simulations, which is a hard task given the above issues.

The above obstructions advocate for collaborative strategies combining physics-driven and data-driven approaches such as the one that we develop in this paper. A similar idea has been explored in [11], but for forecasting temporal series of pollutant, instead of performing state estimation on a large spatial domain.

It should be noted that the limited number of reliable measurements will still pose problems for validating and assessing the quality of each model, which is a crucial part in collaborative strategies. We will mitigate this defect by operating multiple leave-one out cross validations, which preserve as much data as possible for the training part of each model, while testing them on many instances.

1.3 Contributions and layout of the paper

Our main contributions are:

1. the construction of numerous physics-based and data-driven models for state estimation;
2. the construction of a very general ensemble super-learning method combining the above models;

3. its application to the task of recovering urban pollution maps at a city scale, together with a comparison of its constitutive submodels;
4. and the development of a routine extracting car emissions in each street from traffic maps. Moreover, we have created a dataset comprised of processed traffic data from Google Maps screenshots, which can be used for future research.

In our numerical experiments, we work with the inner city of Paris, which covers a surface of about 140 km². The pollutant we consider is NO₂, which is monitored for its respiratory effects, while being mainly produced by vehicle emissions. We use concentration measurements from Airparif sensors¹, and real-time traffic data from Google Maps². Compared to previous contributions and other existing reconstruction methods (see, e.g., [12, 13]), the use of such online traffic data is rather novel. It gives a rough estimation of the spatial density of street traffic, benefits from a very fine spatial resolution, and can be freely updated as frequently as desired, in contrast to many existing approaches which only use time averages of traffic data.

Another distinctive aspect of our approach is the representation of the city by a graph, where nodes and edges correspond to crossroads and street segments, instead of considering an open subset of \mathbb{R}^2 or \mathbb{R}^3 as the spatial domain. This description immediately includes geometric specificities of the agglomeration under study, such as the orientation of each street or the configuration of each neighborhood. It is a natural framework for taking into account pollutant emissions caused by traffic, which are located on the graph. Moreover, it is in adequacy with our goal of estimating local variations in the concentration of pollutants close to the ground, since the streets are isolated from each other at this height.

Physical models can be solved on such domains thanks to the theory of quantum graphs, that is, metric graphs endowed with a differential operator acting on functions defined on the graph (see [14] for details and references). The metric graph structure leads to the definition of suitable and natural function spaces to pose the problem. Of course, several physical models of different complexity could be considered. In this paper, we work with simple elliptic operators, but the model could be refined by considering, for instance, advection-diffusion operators.

The rest of the paper is organized as follows. In Section 2 we present our guiding numerical example of the Parisian area, and the available data. Section 3 explains the different reconstruction methods we have used for our numerical experiments, including ensemble super-learning methods combining the previous ones. By construction, the super-learner has higher approximation power than each individual model. Section 4 discusses how to theoretically quantify performance and optimality of the numerical algorithms, and why leave-one-out is a good way to estimate this performance in practice. We summarize our numerical experiments and provide some illustrations in Section 5. Finally, Appendix A details the mathematical setting for the problem of pollution state estimation on graphs.

¹We extracted data from the Airparif database, which can be found at <https://data-airparif-asso.opendata.arcgis.com>

²The permission to use Google Maps data for non-profit research is stated here: <https://about.google/brand-resource-center/products-and-services/geo-guidelines/#google-maps>




Figure 1: Cropped Google Map screenshot of Paris and the $m = 13$ available stations in the study: red dots represent the projection of the station locations to the nearest vertex in the graph of streets, while the blue crosses correspond to the exact position of the station.

2 Available data and pre-processing

2.1 Pollution sensors

The main information we use consists of direct measurements of the NO_2 concentration field u at Airparif sensor stations. There are $m = 13$ such stations, which are placed at fixed locations

$$r^{\text{obs}} := \{r_1^{\text{obs}}, \dots, r_m^{\text{obs}}\} \in (\mathbb{R}^2)^m,$$

see Figure 1. Each of the stations provides hourly averages of the concentration of nitrogen dioxide, in $\mu\text{g}/\text{m}^3$, of the form

$$z_i = u(r_i^{\text{obs}}) + \eta_i, \quad i = 1, \dots, m,$$

where η_i is some noise with nominal relative error $|\eta_i|/u(r_i^{\text{obs}}) \leq 15\%$.

Note that, in principle, the equations governing the dispersion of pollutants are time dependent. However, as we only have measurements every hour, we opt for a static model, where the state at a given time is computed based on the data available at this time only. This essentially amounts to assuming that the emissions vary slowly over time, and that the system reaches an equilibrium state in less than one hour.

2.2 Meteorological conditions

Wind, as well as stratification effects in the atmosphere due to variations in temperature, play a major role in the dispersion of pollutants [15]. Moreover, the chemical equilibrium between NO and NO_2 depends on the cloud cover [16]. Therefore, we collect the temperature $\theta \in \mathbb{R}$ and the wind speed $w \in \mathbb{R}^2$ at every hour from a weather archive³. These two quantities are treated as global, that is, they are assumed to be constant over the spatial domain.

2.3 Traffic

Car traffic is responsible for more than half the emissions of NO_2 in urban environments [17]. There is an increasing number of available sources that give access to traffic data. In our case,

³See <https://www.windguru.cz>. For the wind, we combined the absolute wind speed with the wind direction to obtain a vector in \mathbb{R}^2 .

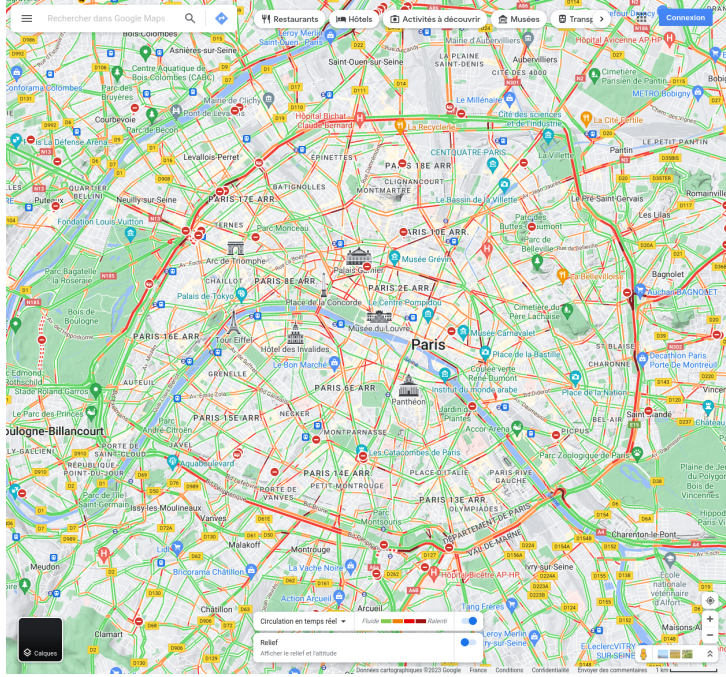


Figure 2: Raw data from Google Maps: the image contains the city with its main landmarks, and some streets are highlighted with one of the four colors corresponding to traffic.

we work with traffic information extracted from Google Maps. We have designed a script using the Python library Selenium to automatically take screenshots of Paris every 15 minutes over an area of 1253×1253 pixels with zoom level 13. An example of resulting *raw image* can be seen in Figure 2. Note that city landmarks could not be removed before taking the screenshot, nor even by subtracting a background image, since each screenshot has slight color variations, rendering this approach impractical. Another issue is the absence of traffic data in the smallest streets of the city. In addition, linking it to the pollution field requires some calibration.

On the one hand, this kind of information is very rich because of its availability in real time, and its spatial coverage of the whole city at high resolution. On the other hand, it is very partial: it comes in the form of four colors, each representing a certain traffic intensity, which gives a qualitative estimate of the number of cars in each street: a street marked in red is, for instance, more congested than one marked in green. One main goal in our work is to examine the potential of incorporating such low-quality information for state estimation tasks.

2.4 Graph of Paris

In order to locate our sources and to express the spatial dependence of a state, we consider a graph domain. We use a metric graph $G = (V, E)$ provided by Open Street Maps, together with the Python library osmnx. For the mathematical definition of metric graphs and their associated function spaces, we refer to Appendix A. The graph G covers the whole inner ring of the city, as shown in Figure 3. The full graph has $|V| = 12963$ vertices and $|E| = 25476$ edges, but we restrict it to the biggest connected component of the subgraph that remains after filtering out all the edges which have never been colored with traffic information. After this operation, our actual graph has $|V| = 10116$ vertices and $|E| = 18713$ edges. The street network is relatively dense, most nodes having 3 to 6 edges.

The vertices $v \in V$ come with precise geographical coordinates. In the following, we assume that the graph is embedded in the two-dimensional plane, and do not take altitude into account.

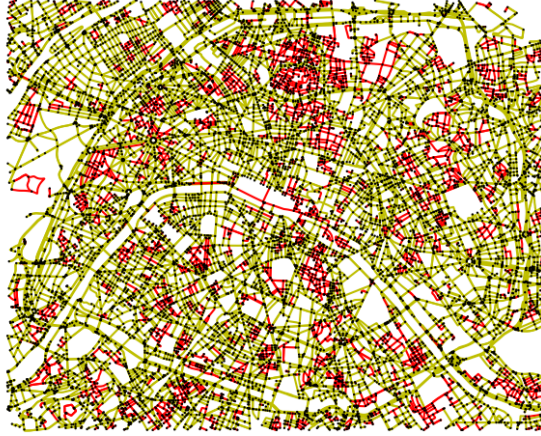


Figure 3: The metric graph downloaded from Open Street Maps, with the edges that never had Google Traffic activation in red, and the edges remaining after filtration in yellow

Each edge $e \in E$ is a street or a portion of it, and we have access to its length ℓ_e as well as its shape, number of lanes and speed limit. The information is so detailed that the streets are represented by paths that are not necessarily straight lines. However, in the following, we will work under this slightly simplified setting.

The location of the sensor stations does not exactly match with vertices of the graph. We therefore project their position r^{obs} onto the nearest vertex, which yields observational nodes

$$v_i^{\text{obs}} := \underset{v \in V}{\operatorname{argmin}} |r_i^{\text{obs}} - v|, \quad i = 1, \dots, m,$$

As Figure 1 illustrates, the projected locations are very close to the exact locations, with a maximal discrepancy of 165m, to be compared with the width of the domain, of about 12km. As a consequence, we will assume that the observations z_i correspond to the values $u(v_i^{\text{obs}})$, up to a slight increase in the measurement errors $|\eta_i|$.

2.5 Pre-processing of traffic data

We also map the traffic information onto pollutant emissions on the graph edges, by implementing the following pipeline:

- **Cropping:** Starting from a raw image like Figure 2, we first crop it to the shape 800×1000 , in order to eliminate toolbars and adapt it to the size of the graph. The background of Figure 1 is obtained by the same procedure.
- **Traffic colors extraction:** The colors associated with the four levels of traffic

$$\text{colors} := \{\text{green}, \text{orange}, \text{red}, \text{dark-red}\}$$

are easily identified⁴. They seem to be used exclusively for that purpose, hence it suffices to extract the pixels having one of these colors.

- **Projection on graph edges:** These pixels, once expressed in their geographical coordinates, almost perfectly overlap the metric graph from Open Street Maps. For each edge $e \in E$ and each color $c \in \text{colors}$, we count the number p_c^e of pixels of color c that are

⁴The RGB value of each color is given by: **green** = (99, 214, 104), **orange** = (255, 151, 77), **red** = (242, 60, 50), **dark-red** = (129, 31, 31)

closest to edge e . Note that the traffic color might change along an edge, in which case we give up on some local information by only considering the total traffic on the edge.

- **Edge normalization:** We then transform these pixel counts into proportions of traffic colors on each edge. As the edges may remain blank at times where there is no traffic, we take as a normalizing constant the maximal amount of pixels encountered over all times T for which we collect traffic data:

$$q_c^e(t) = \frac{p_c^e(t)}{\max_{t' \in T} \sum_{c' \in \text{colors}} p_{c'}^e(t')}, \quad t \in T.$$

In this way, $q_c^e(t) \in [0, 1]$ indicates the proportion of edge e colored with c at time t , but remains null if no traffic is reported.

- **Hourly averaging:** As the pollution information is only available every hour, we take the average of the four values of q_c^e encountered every fifteen minutes, which we still denote q_c^e in the sequel.
- **Projection on graph nodes:** In our models, it is in fact simpler to localize emissions on the nodes of the graph. For this reason, we calculate the density of each traffic color c around a vertex $v \in V$ as a weighted average on its neighboring edges $E(v)$

$$q_c^v(t) = \frac{\sum_{e \in E(v)} a_e q_c^e(t)}{\sum_{e \in E(v)} a_e},$$

where a_e stands for the area of the road associated to edge e , given by the product of its length ℓ_e with the number of lanes.

2.6 Summary

While the history of sensor and weather data can be found on archives, our script for capturing traffic images only runs in real time, since Google Traffic only provides current information. We collected all types of data on an hourly basis for a period of time comprised between December 9, 2022 and March 19, 2023. After removing time stamps for which some data was missing, we end up with a set of acquisition times T , of cardinality $|T| = 1712$, which we divide into a set T_{train} of 1338 training times, and a set T_{test} of 374 testing times.

In the end, given the graph $G = (V, E)$, the available information at any time $t \in T$ is of the form

$$x = (v^{obs}, z, \theta, w, (q_c^v)_{c,v}) \in \mathcal{X} = V^m \times \mathbb{R}^m \times \mathbb{R} \times \mathbb{R}^2 \times \mathbb{R}^{4|V|}. \quad (2.1)$$

In the next section, we present various models to estimate the pollution field from this data, using either statistical inference, linear mappings based on expert knowledge, or neural networks.

3 Reconstruction methods

We have implemented several methods of state estimation by leveraging the different information sources. Our methods give reconstructions on the metric graph G , that is, we consider mappings $A : \mathcal{X} \rightarrow U$ where U is a space of functions defined on G . Typical examples are $U = \mathcal{C}(G)$, $L^2(G)$ or $H^1(G)$, as defined in Appendix A. As our main interest is in assessing the effect of incorporating indirect information like the real-time traffic data, we first consider models that take only a portion of $x \in \mathcal{X}$ as input.

3.1 Spatial average

If we give up on all the spatially-dependent data v^{obs} and $(q_c^e)_{c,e}$, the reconstruction is necessarily constant over the whole domain \mathbf{G} , which yields no better choice than the average of the observed concentration values

$$A_{\text{avg}}(x)(r) = \bar{z} = \frac{1}{m} \sum_{i=1}^m z_i, \quad \forall r \in \mathbf{G}.$$

This extremely simple reconstruction will serve as our baseline to compare more sophisticated reconstructions. In the sequel, we will add the spatially-dependent data and view the other models as corrections to the spatial average above. This will result in spatially unbiased estimators, provided that the locations of the sensors are representative of the whole pollution field. More precisely, assuming that the stations are randomly drawn according to the uniform probability distribution μ on \mathbf{G} , the expectation over z_1, \dots, z_m of the spatially-averaged error is

$$\mathbb{E}_z \left(\int_{\mathbf{G}} (u(r) - A_{\text{avg}}(x)(r)) d\mu(r) \right) = \int_{\mathbf{G}} u d\mu - \mathbb{E}_z \left(\frac{1}{m} \sum_{i=1}^m z_i \right) = 0,$$

and this remains true when adding to A_{avg} a correction of vanishing expectation.

3.2 Best unbiased linear estimator

If we only want to estimate a missing measurement z_i at a given station $i \in \{1, \dots, m\}$, we may also use statistical information stemming from the history (z_i^t) of the station at previous times t , as well as the observations from other stations $j \neq i$ in the present and the past, denoted respectively z_j and (z_j^t) . For T_{train} the set of training times, define the empirical average

$$\langle z_i \rangle := \frac{1}{|T_{\text{train}}|} \sum_{t \in T_{\text{train}}} z_i^t$$

and empirical covariance matrix $K \in \mathbb{R}^{m \times m}$ with entries

$$K_{i,j} := \left\langle (z_i - \langle z_i \rangle)(z_j - \langle z_j \rangle) \right\rangle = \langle z_i z_j \rangle - \langle z_i \rangle \langle z_j \rangle.$$

Any unbiased linear estimator \tilde{z}_i of z_i is of the form

$$\tilde{z}_i = \langle z_i \rangle + \sum_{j \neq i} c_j (z_j - \langle z_j \rangle),$$

for some coefficients $(c_j)_{j \neq i}$. Let $c \in \mathbb{R}^m$ be the vector with coordinates c_j for $j \neq i$ and $c_i = -1$. Then the best linear unbiased estimator (BLUE) is obtained by optimizing the averaged squared error

$$\underset{(c_j)_{j \neq i}}{\operatorname{argmin}} \left\langle (\tilde{z}_i - z_i)^2 \right\rangle = \underset{(c_j)_{j \neq i}}{\operatorname{argmin}} c^\top K c = [(K_{j,k})_{j,k \neq i}]^{-1} (K_{j,i})_{j \neq i},$$

where $(K_{j,k})_{j,k \neq i}$ and $(K_{j,i})_{j \neq i}$ are seen as a matrix in $\mathbb{R}^{(m-1) \times (m-1)}$ and a vector in \mathbb{R}^{m-1} .

If the set of training times T_{train} is large enough, we expect an ergodicity property of the form $\langle z_i \rangle \approx \mathbb{E}(z_i)$ to hold. For this reason, BLUE should be a near minimizer of the expected squared error, given the available data. Therefore, in the numerical experiments, we will evaluate the different methods A by comparing $A(x \setminus \{z_i\})(z_i)$ and z_i , and the error $|\tilde{z}_i - z_i|^2$ will act as an optimality benchmark.

It should be emphasized that BLUE itself is not a valid reconstruction method, since it requires statistical information which is accessible only at the locations of the stations v_i^{obs} . Hence this estimator cannot be computed at any point $r \in \mathbf{G}$ of the graph domain.

3.3 Kriging

In order to transform BLUE into a reconstruction method, one needs to propose a surrogate for the correlation between any two points in the graph. Moreover, as we don't know the average pollution at all points of the graph, we proceed without subtracting spatial averages $\langle z_i \rangle$ in this subsection, in contrast to the previous one. Therefore, we consider the Gram matrix of normalized second-order moments

$$G_{i,j} = \frac{\langle z_i z_j \rangle}{\sqrt{\langle z_i^2 \rangle \langle z_j^2 \rangle}}.$$

Taking the positions \mathbf{v}^{obs} of the stations into account, we observe that each entry $G_{i,j}$ partly depends on the distance $|\mathbf{v}_i^{\text{obs}} - \mathbf{v}_j^{\text{obs}}|$ between the stations, see Figure 4. A typical choice of approximant is the Gaussian kernel

$$G_{i,j} \approx \hat{G}_{i,j} := C \exp\left(-\frac{|\mathbf{v}_i^{\text{obs}} - \mathbf{v}_j^{\text{obs}}|^2}{2\sigma^2}\right) + (1 - C)\delta_{i,j},$$

with parameter values $C = 0.968$ and $\sigma = 33.4\text{km}$ obtained by fitting the station data in our case.

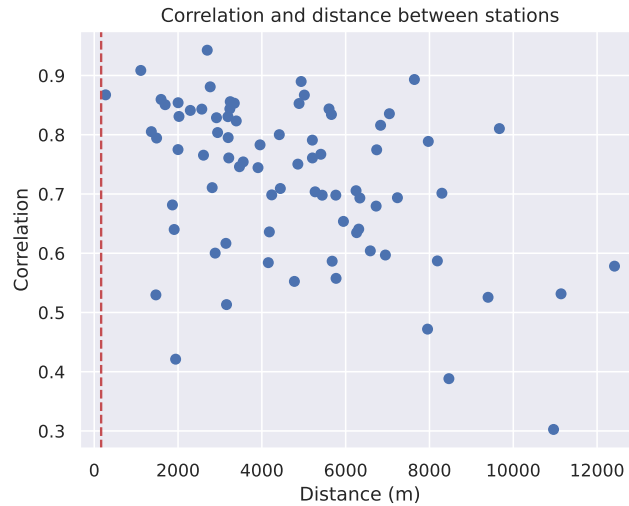


Figure 4: Correlation between stations as a function of the distance. The vertical slashed red line marks the maximal separation between vertex and station (165m) which still lays in the zone of high correlation.

Remark 3.1. The fact that $C < 1$ can be interpreted as the presence of random noise η_i on the measurements $z_i = u(\mathbf{v}_i^{\text{obs}}) + \eta_i$. As a safety check, one may notice that the average relative error

$$\frac{\langle \eta_i^2 \rangle}{\langle u(\mathbf{v}_i^{\text{obs}})^2 \rangle} = \frac{\langle z_i^2 \rangle}{\langle u(\mathbf{v}_i^{\text{obs}})^2 \rangle} - 1 \approx \frac{1}{C} - 1 = 3.31\%$$

is effectively much smaller than the uniform error guarantee $\|\eta_i / |u(\mathbf{v}_i^{\text{obs}})|\|_{L^\infty} \leq 15\%$ that we discussed in Section 2.1. Adding the matrix $(1 - C)I$ ensures that \hat{G} has ones on its diagonal, as expected of a correlation matrix, and regularizes the system, by making the inversion of \hat{G} stable. More practically, the reconstructed value $A(x)(\mathbf{v}_i^{\text{obs}})$ will not be exactly z_i , but rather an average of the measurements close to $\mathbf{v}_i^{\text{obs}}$.

Let $r \in \mathbb{G}$, we again examine a linear model

$$A_{\text{krig}}(x)(r) = \sum_{i=1}^m c_i^r z_i,$$

with coefficients $c^r \in \mathbb{R}^m$ to be determined. This estimator is unbiased if and only if $\sum_{i=1}^m c_i = 1$, and in that case we can write it as a correction to the temporal or spatial average

$$A_{\text{krig}}(x)(r) = \langle A_{\text{krig}}(x)(r) \rangle + \sum_{i=1}^m c_i^r (z_i - \langle z_i \rangle) = \bar{z} + \sum_{i=1}^m \left(c_i^r - \frac{1}{m} \right) (z_i - \bar{z}).$$

By analogy with BLUE, we thus define the weights as the renormalized solution of a system of correlation equation

$$c^r = \frac{\hat{c}^r}{\sum_{i=1}^m \hat{c}_i^r}, \quad \hat{c}^r = \hat{G}^{-1} g^r, \quad g^r = C \exp\left(-\frac{|v_i^{\text{obs}} - r|^2}{2\sigma^2}\right).$$

Although there are no optimality guarantees, we expect kriging to have intermediate performance when compared to the spatial average baseline, and to the ideal BLUE reconstruction. However, due to the important spacing between peripheral stations, we only observe an improvement in the central part of Paris. The insufficient density of pollution measurements calls for models involving other sources of information, such as traffic data. This is the objective of the next two subsections.

3.4 Source model

The simplest way to incorporate traffic data consists in using only local values q_c^e for estimating the pollution on an edge $e \in \mathbb{E}$, or q_c^v for a node $v \in \mathbb{V}$. As we projected the station locations on \mathbb{V} , we focus on the latter case here. We call such a method a source model, since it directly maps the sources of emission to pollution values.

We opt for a linear model acting as a correction on the spatial average baseline:

$$A_{\text{src}}(x)(v) = \bar{z} + \sum_{c \in \text{colors}} \alpha_c (q_c^v - \bar{q}), \quad (3.1)$$

where we subtracted the spatial average of traffic \bar{q} for unbiasedness. The vector of coefficients $\alpha \in \mathbb{R}^4$ is found by solving a LASSO problem

$$\min_{\alpha \in \mathbb{R}^4} \sum_{t \in T_{\text{train}}} \sum_{i=1}^m |z_i^t - A_{\text{src}}(x(t))(v_i^{\text{obs}})|^2 + \lambda \|\alpha\|_1,$$

and we perform a cross-validation to estimate the optimal parameter λ , in order to prevent overfitting.

Alternatively, we can also write nonlinear variants, of the form

$$A_{\text{src}}(x)(v) = \bar{z} + \mathcal{T}_\alpha((q_c^v), \theta, w), \quad (3.2)$$

which may take into account other sources of information like temperature θ and wind w . Here, $\mathcal{T} : \mathbb{R}^{\# \alpha} \times \mathbb{R}^4 \times \mathbb{R} \times \mathbb{R}^2 \rightarrow \mathbb{R}$ can be a *polynomial* combination of the inputs $((q_c^v), \theta, w)$, a *neural network*, or any other nonlinear mapping. The set of parameters α is no longer associated to the four traffic colors, but still needs to be learned via a LASSO regression.

All such models rely on the assumption that pollution depends on its sources in a very localized manner. However, the traffic charts $(q_c^v)_{v \in \mathbb{V}}$ exhibit sharp variations from one node to its neighbors, which incites to smooth the emissions before applying the above methods. This is the purpose of the following section, which attempts to model such a diffusive behavior.

3.5 Physical modeling

An important inspiration for reconstruction methods comes from the physical modeling of pollution dispersion. In our setting, we resort to building a quantum graph, that is, we endow our metric graph G with a differential operator acting on functions from functional spaces such as $L^2(\mathsf{G})$ or $H^1(\mathsf{G})$, as defined in Appendix A (we also refer to [14] for more details and references). Our approach can be summarized as follows:

Elliptic equation: One can first model pollution with a time-independent elliptic equation, by assuming that all time-dependent parameters have sufficiently slow variations, here over the course of an hour, for the pollution field to reach a steady state. For any point $r \in \mathsf{G}$, the pollutant concentration is modelled by a function $u : \mathsf{G} \rightarrow \mathbb{R}$ solution to the diffusion-reaction equation

$$\mathcal{P}(u) := -\frac{d}{dr} \cdot \left(a(r) \frac{d}{dr} u(t, r) \right) + h(r)u(r) = q(r), \quad r \in \mathsf{G} \quad (3.3)$$

which we choose to complement with “Newmann-Kirchoff” conditions on the vertices, that is,

$$\sum_{e \in \mathsf{E}(\mathbf{v})} \left. \frac{du}{dr} \right|_e (\mathbf{v}) = 0, \quad \mathbf{v} \in \mathsf{V}, \quad (3.4)$$

expressing the conservation of the quantity of pollutant at every crossroad $\mathbf{v} \in \mathsf{V}$. Here, $\mathsf{E}(\mathbf{v})$ denotes the edges having \mathbf{v} as an endpoint, and the derivatives are assumed to be taken in the directions away from the vertex.

In equation (3.3), the function $a \in L^\infty(\mathsf{G})$ is an effective diffusion coefficient, which takes into account turbulent dissipative effects. The absorption coefficient $h \in L^\infty(\mathsf{G})$ models the leakage of pollutants from the streets to the higher atmosphere, as well as chemical reaction, in particular between NO_2 and other nitrogen oxides, which are not measured by the sensors. Lastly, the source term $q \in L^2(\mathsf{G})$ models all possible emissions of pollutant, which in the case of Paris essentially come from traffic, local heating, and industrial and urban activity outside the city. As we only have access to local traffic data, we assume that the other sources are spatially constant, and average them out by solving $\mathcal{P}(u - \bar{u}) = q - \bar{q}$, where \bar{u} is the spatial average of the pollutant concentration, estimated by $A_{\text{avg}}(x) = \bar{z}$, and $q - \bar{q}$ corresponds to the variations of traffic around its spatial average, computed through the procedure from Sections 2.5 and 3.4.

Remark 3.2. Equation (3.4) has a similar effect as Newmann conditions at the borders of the spatial region under consideration, here the rectangle contour of Figure 1. Therefore the only exchanges with the exterior of this region are contained in the source term q . As this no-flux condition only give a very rough approximation of the solution close to the border, in the numerical experiments of Section 5, we will concentrate on the accurate prediction of the pollution in the central part of the city.

Variational formulation: The operator $\mathcal{P}(u)$ in (3.3) is defined for functions $u \in H^2(\mathsf{G})$, but the equation can be stated in a weak form, which only requires that $u \in H^1(\mathsf{G})$. Multiplying (3.3) by a sufficiently smooth test function $v \in H^1(\mathsf{G})$, and using the Kirchoff-Neumann boundary conditions, it follows that the corresponding weak formulation of the problem is to find $u \in H^1(\mathsf{G})$ such that

$$\mathfrak{b}(u, v) = \mathfrak{f}(v), \quad v \in H^1(\mathsf{G}) \quad (3.5)$$

where \mathfrak{b} is the symmetric bilinear form defined as

$$\begin{aligned} H^1(\mathsf{G})^2 &\rightarrow \mathbb{R} \\ \mathfrak{b} : (u, v) &\mapsto \sum_{e \in \mathsf{E}} \left\{ \int_e a(r) \frac{du}{dr}(r) \frac{dv}{dr}(r) dr + \int_e h(r)u(r)v(r) dr \right\} \end{aligned}$$

and $\mathbf{f} : v \in H^1(\mathbf{G}) \mapsto \sum_{\mathbf{e} \in \mathbf{E}} \int_{\mathbf{e}} q(r)v(r)dr$ is a continuous linear form.

Assuming that $a(r) \geq a_0 > 0$ and $h(r) \geq h_0 > 0$ for $r \in \mathbf{G}$ a.e., we see that \mathbf{b} is continuous and coercive in $H^1(\mathbf{G})$ with coercivity constant $\min(a_0, h_0)$, and continuity constant $\max(\|a\|_{L^\infty(\mathbf{G})}, \|h\|_{L^\infty(\mathbf{G})})$. By the Lax-Milgram theorem, problem (3.5) admits a unique solution $u \in H^1(\mathbf{G})$.

Discretization: In our numerical tests, we discretize the equation with \mathbb{P}_1 finite elements, that is, continuous functions whose restriction to any edge is affine. We describe below the main guidelines, and refer to [18] for further details and a complete analysis.

We define the set of hat functions $\{\varphi_{\mathbf{v}}\}_{\mathbf{v} \in \mathbf{V}}$ by $\varphi_{\mathbf{v}}(\mathbf{v}') = \delta_{\mathbf{v}, \mathbf{v}'}$ for any vertices $\mathbf{v}, \mathbf{v}' \in \mathbf{V}$, and

$$\forall x_{\mathbf{e}} \in [0, \ell_{\mathbf{e}}], \quad \varphi_{\mathbf{v}}(x_{\mathbf{e}}) = \begin{cases} 1 - \frac{x_{\mathbf{e}}}{\ell_{\mathbf{e}}}, & \text{if } \mathbf{e} \in \mathbf{E}(\mathbf{v}), \\ 0, & \text{if } \mathbf{e} \notin \mathbf{E}(\mathbf{v}), \end{cases}$$

for any edge $\mathbf{e} \in \mathbf{E}$. Fixing our finite element space $\mathbb{P}_1 = \text{span}\{\varphi_{\mathbf{v}}\}_{\mathbf{v} \in \mathbf{V}} \subset H^1(\mathbf{G})$, we search for the Galerkin solution $\hat{u} = \sum_{\mathbf{v} \in \mathbf{V}} c_{\mathbf{v}} \varphi_{\mathbf{v}} \in \mathbb{P}_1$ such that

$$\mathbf{b}(\hat{u}, \hat{v}) = \mathbf{f}(\hat{v}), \quad \hat{v} \in \mathbb{P}_1.$$

Gathering the expansion coefficients of the solution in the vector $\mathbf{c} = \{c_{\mathbf{v}}\}_{\mathbf{v} \in \mathbf{V}}$, we obtain the linear system of equations

$$\mathbb{B} \mathbf{c} = \mathbf{f} \tag{3.6}$$

with $\mathbb{B} = (\mathbf{b}(\varphi_{\mathbf{v}}, \varphi_{\mathbf{v}'}))_{\mathbf{v}, \mathbf{v}' \in \mathbf{V}}$ and $\mathbf{f} = (\mathbf{f}(\varphi_{\mathbf{v}}))_{\mathbf{v} \in \mathbf{V}}$. Again by Lax-Milgram theory, this system is invertible, which allows to compute the solution \hat{u} .

Reduced models: Unfortunately, solving equation (3.6) is expensive, given the size $|\mathbf{V}| \approx 10^4$ of the graph, so we cannot afford to find \hat{u} at each time step. In order to mitigate the computational cost, we rely on model order reduction techniques, which have received much attention in the context of parametrized elliptic PDEs [19, 20, 21, 22, 23, 24]. Here, the parameters would be the diffusion a , the reaction h , and the right-hand side q . We consider three reconstructions methods.

1. **Eigenstates of the graph Laplacian:** One option consists in taking as a reduced model the subspace $V_n \subset H^1(\mathbf{G})$ spanned by the n first eigenfunctions of the Laplacian in \mathbb{P}_1 . As this operator is self-adjoint and coercive, it admits a spectral decomposition with positive eigenvalues, and the coefficients of the eigenstates in the basis $\{\varphi_{\mathbf{v}}\}_{\mathbf{v} \in \mathbf{V}}$ are the eigenvectors of \mathbb{B} . Assuming that the diffusion and reaction coefficients a and h are constants calibrated in a pre-processing phase, we define the reconstruction mapping $A : \mathcal{X} \rightarrow H^1(\mathbf{G})$ by taking $\hat{u} = A(q)$ the Galerkin projection of u onto V_n , that is, by searching $\hat{u} \in V_n$ solution to

$$\mathbf{b}(\hat{u}, \hat{v}) = \mathbf{f}(\hat{v}), \quad \hat{v} \in V_n,$$

which is simply a diagonal system in the eigenstate basis. We then plug \hat{u} instead of q in equation (3.1) or (3.2), and learn the coefficients associated to each color, or the more general parameters α .

2. **Principal components of traffic data:** Starting with the whole history of traffic data $(q(t))_{t \in T} \in \mathbb{R}^{|T| \times |\mathbf{M}|}$, we can also perform a singular value decomposition to find the n first modes q_1, \dots, q_n , compute the solutions to $\mathcal{P}(u_k) = q_k$, and assemble them in a reduced space $V_n = \text{span}\{u_1, \dots, u_n\}$. In this way, we expect V_n to better capture physical properties of the pollution field, such as strong correlations along a large avenue.

As full-order solves remain costly, we resort to a convolution with a gaussian kernel:

$$(u_k)_c^v = \frac{\sum_{v' \in V} e^{-d_{vv'}^2/2\delta^2} (q_k)_c^{v'} \phi_{v'}}{\sum_{v' \in V} e^{-d_{vv'}^2/2\delta^2}}$$

where $d_{vv'} = |v - v'|$ is the distance in \mathbb{R}^2 (which is equivalent, up to constants, to the distance on the graph). We set $\delta = 400$ m, after observing that pollution data is optimally correlated to regularized traffic information for δ close to this value.

In our experiments, we perform the smoothed projections $q \mapsto \hat{u} = \sum_{k=1}^n \langle q, q_k \rangle u_k$ into a different reduced space for each of the four traffic colors. After this operation, we can apply any of the strategies described in section 3.4 to \hat{u} instead of q .

These two methods regularize the traffic data, but they do not exploit the information from the pollution sensors, apart from the average value \bar{z} . In order to assimilate data of both types, it is possible to use a combined least-squares fit of the form

$$A(x) = \operatorname{argmin}_{\hat{v} \in V_n} \|z - \hat{v}(\mathbf{v}^{\text{obs}})\|_2^2 + \lambda' \|q - \hat{q}\|_{\ell^2(V)}^2,$$

where $\lambda' > 0$ balances the contributions of z and q , and $\hat{q} = \mathcal{P}(\hat{v}) = \sum_{k=1}^n \hat{c}_k q_k$ for the coefficients $\hat{c} \in \mathbb{R}^n$ such that $\hat{v} = \sum_{k=1}^n \hat{c}_k u_k$. However this approach did not perform well in practice, so we did not include it in the numerical experiments.

In the last method, if $m \geq n$ and λ tends to 0, the prediction \hat{u} does a least squares fit of u at the available measurement points $\mathbf{v}_i^{\text{obs}}$. In general, it is possible to enforce $\hat{u}(\mathbf{v}^{\text{obs}}) = u(\mathbf{v}^{\text{obs}})$ by applying a correction to the prediction. This post-process, called *Parameterized Background Data-Weak* method, was originally introduced in [25] and has been analyzed and extended in a series of papers such as [26, 27, 28, 29]. The whole approach has found numerous applications, including pollution dispersion [30].

It would of course be possible to gain in accuracy, by considering more refined equations for pollution dispersion, which capture additional physical properties, and thus by encoding these properties into the reduced space V_n . One could for instance think of advection by wind, vertical fluxes or stratification of the atmosphere depending on the temperatures, changes in the chemical equilibrium between NO and NO₂ caused by cloud coverage and precipitations [16], as well as local turbulent effects near the sensor stations. Note that, if nonlinear equations are involved, V_n can be a nonlinear approximation space defined through a chart of n parameters, and approximation guarantees are more difficult to obtain [29].

3.6 Super-Learning as a collaborative approach

To gain in accuracy over each individual model, one can combine a set of p available mappings A_1, \dots, A_p coming from the previous methods, and build a super-learner

$$\begin{aligned} \mathcal{F}(\mathcal{X}, U)^p &\longrightarrow \mathcal{F}(\mathcal{X}, U) \\ \mathcal{S} : (A_1, \dots, A_p) &\longmapsto \mathcal{S}(A_1, \dots, A_p), \end{aligned}$$

where $\mathcal{F}(\mathcal{X}, U)$ denotes the set of functions from \mathcal{X} to U . The most simple merger, usually called aggregator in statistics, amounts to taking a linear combination

$$\mathcal{S}_\omega(A_1, \dots, A_p) = \sum_{i=1}^p \omega_i A_i,$$

for some weights $\omega = (\omega_1, \dots, \omega_p)$ expressing the confidence in each individual model.

More sophisticated strategies involve nonlinear combinations and compositional structure. One could think of using a first model to obtain a rough estimation, and compose it with a second model performing refinements based on its output. This is already an underlying idea in our constructions, where we start with the spatial average, and add spatially-dependent corrections. The physical models involve one more compositional step, since they are of the form $A_{\text{src}} \circ \hat{u}(q)$.

Neural networks constitute another prominent example of nonlinear super-learners: one could treat $A_1(x), \dots, A_p(x)$ as inputs, and train the parameters ω by minimizing an empirical loss. We would like to emphasize here that properly training the super-learner requires to implement a nested leave-one-out strategy: one should first train each parametrized submodel by leave-one-out, and then optimize the neural network with another leave-one-out step, in order to avoid overfitting. As a consequence, at least two observation points are removed from the training set of the submodels, which may cause a loss of accuracy, especially when the number m of observations is small.

In our application, the neural network super-learner performed slightly worse than its linear counterpart, which should come as no surprise in view of the above observation.

4 Reconstruction benchmarks and Leave-One-Out

There are several ways to quantify the quality of a reconstruction map $A : \mathcal{X} \rightarrow U$. Ideally, given a state $u \in U$ and the associated observations $x \in \mathcal{X}$, one would like to find A such that the error $\|u - A(x)\|_U$ is as small as possible. Assuming that (u, x) is a random variable with distribution $\pi \in \text{Prob}(U \times \mathcal{X})$, we define the performance of A as the L^2 norm of the error

$$e(A)^2 := \int_{U \times \mathcal{X}} \|u - A(x)\|_U^2 d\pi(u, x),$$

which acts as a good compromise between the worst-case error and the average error. In addition, although the state u has in principle H^2 regularity, we assess the spatial error also in L^2 , that is, we take $U = L^2(\mathbf{G})$. Unfortunately, finding A minimizing $e(A)$ is out of reach for several reasons. First, we don't know the distribution π , nor even its support, which is the set of all possible states and observations. Second, given $u \in U$, we cannot evaluate $u(r)$ at any point $r \in \mathbf{G}$, making the computation of $\|u - A(x_u)\|_U$ intractable.

Concerning the first issue, as we have access to hourly data on a large period of time, we can replace the integral over π by an empirical average

$$e(A)^2 \approx \frac{1}{T_{\text{test}}} \sum_{t \in T_{\text{test}}} \|u(t) - A(x)(t)\|_U^2$$

over the set T_{test} of 374 test times. Assuming that these states are independent, this approximation induces an error of order

$$\frac{\mathbb{E}(\|u\|_U^2)^{1/2}}{\sqrt{|T_{\text{test}}|}} \approx \frac{41}{19.3} \approx 2.1 \mu\text{g}/\text{m}^3,$$

which is totally acceptable in view of the noise level on the measurements.

The second obstacle is more tricky, because we only know u at a very limited number m of fixed positions, and because these observations are also needed for constructing A . Ignoring the last issue leads to a systematic underestimation of $e(A)$, as we detail below.

Assume that the observation points v_i^{obs} are distributed uniformly at random on \mathbf{G} , define the discrete semi-norm

$$\|u\|_m^2 = \frac{1}{m} \sum_{i=1}^m |u(v_i^{\text{obs}})|^2$$

corresponding to an empirical version of $\|u\|_U^2$, and consider map A solution to

$$A: \mathcal{X} \rightarrow V_n \text{ linear} \quad \min \frac{1}{T_{\text{train}}} \sum_{t \in T_{\text{train}}} \|u - A(x)\|_m^2$$

for some linear space $V_n \subset U$ of dimension n . This setting is valid for most of our methods, with V_n the set of constant functions in the case of A_{avg} (of dimension $n = 1$), but also $V_n = \text{span}\{r \mapsto c_i^r\}_{1 \leq i \leq m}$ in the case of A_{krig} (of dimension $n = m$), and V_n the reduced basis in the methods based on physical modeling.

Then, for A^* the optimal map taking values in V_n

$$A^* = \underset{A': \mathcal{X} \rightarrow V_n \text{ linear}}{\text{argmin}} \quad \mathbb{E}(\|u - A'(x)\|_U^2) = \mathbb{E}_\pi(u|x),$$

we obtain, by applying Pythagoras theorem both for $\|\cdot\|_U$ and $\|\cdot\|_m$,

$$e(A)^2 = \mathbb{E}(\|u - A(x)\|_U^2) \geq \mathbb{E}(\|u - A^*(x)\|_U^2) = \mathbb{E}(\|u - A^*(x)\|_m^2) \geq \mathbb{E}(\|u - A(x)\|_m^2),$$

where the central equality comes from the assumption that the v_i^{obs} are random. This proves that $\|u - A(x)\|_m^2$ is a biased estimator for $e(A)^2$ as soon as one of the inequalities is strict, that is, as soon as $A(x) \neq A^*(x)$. Note that separating the training and test data by splitting the set of time indices T is not sufficient, since the algorithm will then fit its prediction to the station locations, without generalization guarantees to the rest of the domain G .

As a consequence, we must separate the stations into a training set and test points. In order to compute an unbiased estimator of $e(A)^2$ with minimal variance, while keeping the maximal number of stations in the training set, we proceed to leave-one-out cross-validation. This procedure is very standard and has been used in other works on pollution reconstruction, such as [10]. For $1 \leq i \leq m$, denote

$$\|u\|_{m \setminus i}^2 = \frac{1}{m-1} \sum_{j \neq i} |u(v_j^{\text{obs}})|^2 \quad \text{and} \quad A_i = \underset{A: \mathcal{X} \rightarrow V_n \text{ linear}}{\text{argmin}} \frac{1}{T_{\text{train}}} \sum_{t \in T_{\text{train}}} \|u - A(x)\|_{m \setminus i}^2.$$

Assuming that the station locations are independent random variables, the cross-validation estimator of the error

$$e_{\text{CV}}(A)^2 := \frac{1}{T_{\text{test}}} \sum_{t \in T_{\text{test}}} \frac{1}{m} \sum_{i=1}^m |u(v_i^{\text{obs}}) - A_i(v_i^{\text{obs}})|^2$$

is unbiased, since

$$\mathbb{E}(e_{\text{CV}}(A)^2) = \mathbb{E}(|u(v_i^{\text{obs}}) - A_i(v_i^{\text{obs}})|^2) = \mathbb{E}(\|u - A_i(x)\|_U^2) = e(A),$$

with the difference that x contains only $m - 1$ direct evaluations of u this time.

5 Numerical results

We have implemented and tested numerous variants and combinations of the models from section 3. This was done thanks to a Python code we have developed, which can be found at <https://github.com/agussomacal/CityPollutionModeling>. The interested user could add its own models for further testing. In this section, we summarize the most important results that emerge from our tests. We report on the performance of the following reconstruction strategies:

- **Spatial average:** We take a simple spatial average, as in section 3.1. The resulting error serves as a baseline, which we expect to beat with the other more sophisticated models.
- **BLUE:** As explained in section 3.2, BLUE can be seen as an estimate of the optimal linear reconstruction method. It can be used as a benchmark of the best performance that we can expect of linear methods. Note that, in principle, nonlinear strategies could be more accurate than BLUE. However, we will see in our experiments that none of our methods achieves such accuracy.
- **Kriging:** We apply the kriging method depicted in section 3.3 with an exponential kernel. The parameters σ and C are obtained by fitting an exponential to the correlation between training stations (that is, we do not into account correlations with the station that is set aside for testing) as a function of their distance (see fig. 4).
- **Source:** We apply a linear source model, as described in section 3.4, with temperature θ and wind w as extra regressor variables.
- **Physical-PCA:** We apply the second physical model from section 3.5, using a gaussian kernel to smooth the node traffic data. The four reduced spaces V_n , associated to the four traffic colors, each consist of the first 10 principal components of the corresponding traffic data, as observed in the training set. After the smoothing and projection operations, we assemble the variables θ , w and the q_c^v on each node v into a vector $s = (q_{\text{green}}^v, q_{\text{orange}}^v, q_{\text{red}}^v, q_{\text{dark-red}}^v, \theta, w) \in \mathbb{R}^6$ and build a *polynomial* model of degree 2:

$$\mathcal{T}_\alpha(s) := \sum_{j=1}^6 \alpha_j s_j + \sum_{j,k=1}^6 \alpha_{jk} s_j s_k,$$

where $\alpha \in \mathbb{R}^{42}$ is computed following the lines of section 3.4.

- **Physical-Laplacian:** We apply the first physical model from section 3.5, with the reduced space consisting of the first 5 eigenvectors of the graph laplacian. After projection into the subspace, we take the 4-colour traffic values on each node $(q_c^n)_{\cdot,i}$ and obtain their *degree-3 polynomial* combinations. Finally we apply a *neural network* consisting of two hidden layers of 20 neurons each and a ReLU activation function. The neural network is trained with ADAM optimizer with early stopping to prevent overfitting.
- **Ensemble:** We apply an ensemble model, as described in section 3.6, combining the Kriging method A_{krig} , the Source model A_{src} and the Physical-Laplacian model A_{lapl} . We train each of them separately and compute the following linear combination:

$$A_{\text{ens}}(x)(r) = \omega(r)A_{\text{krig}}(x)(r) + \frac{1 - \omega(r)}{2}A_{\text{src}}(x)(r) + \frac{1 - \omega(r)}{2}A_{\text{lapl}}(x)(r),$$

with a weight function $\omega(r) = \exp(\min_{1 \leq i \leq m} |r - v_i^{\text{obs}}|/\delta)$, where $\delta = 800$ m. Essentially, we favour Kriging when r is close to one of the sensor stations, and average the predictions of models using local or global traffic information otherwise.

In fig. 5, we show the root mean square error (in $\mu\text{g}/\text{m}^3$) for each model's predictions on the test times T_{test} and tested stations i :

$$e_{\text{RMSE}}(A, i) := \left(\frac{1}{T_{\text{test}}} \sum_{t \in T_{\text{test}}} |u(v_i^{\text{obs}}) - A_i(v_i^{\text{obs}})|^2 \right)^{1/2}.$$

Note that the cross-validation error $e_{CV}(A)$ from section 4 is just the ℓ^2 -average of these errors over all stations. For the tests, we only keep the 10 stations located in the interior of Paris. The remaining 3 are set aside because they have a significant proportion of missing values (in average 10% of the data is lacking in each of these stations), and because they lay close to the border of the image, making the surrounding traffic information incomplete.

The shaded blue area is the region corresponding to errors smaller than the reference *BLUE* model. On the opposite side, the shaded red area marks situations in which the prediction is worse than the *spatial average* baseline. The white margin in between indicates the region where we expect feasible improvements.

We first notice that using a linear *source model* already yields reliable improvements with respect to the *spatial average* baseline. It fails, however, in HAUS and OPERA stations due to the absence of traffic information around the former, as the Google Maps symbol for the Paris Opera is drawn over the location of the sensor. This affects the predictions on both stations but most prominently on HAUS. This problem can be alleviated if we average traffic information over a bigger region, as done in *Physical-PCA* thanks to the Gaussian smoothing, at the expense of losing precision on other stations like PA18.

The *Kriging* model manages to produce enhanced predictions in both HAUS and OPERA stations because of their proximity and correspondingly high correlation in pollution values. However, for other stations, especially those further from the center, the performance highly deteriorates.

With the *Physical-Laplacian* model, we get further improvements in 6 stations compared to the linear *source model*, while losing some advantage in the remaining 4. Finally, the *ensemble* method, by combining two traffic models and the *Kriging* method, manages to exploit the advantages of each in a pretty decent way. It yields predictions that beat or equal the *spatial average* baseline on all stations, and that reach the best average error among all our tested methods. In fig. 6, we show an example of pollution maps generated with this last model.

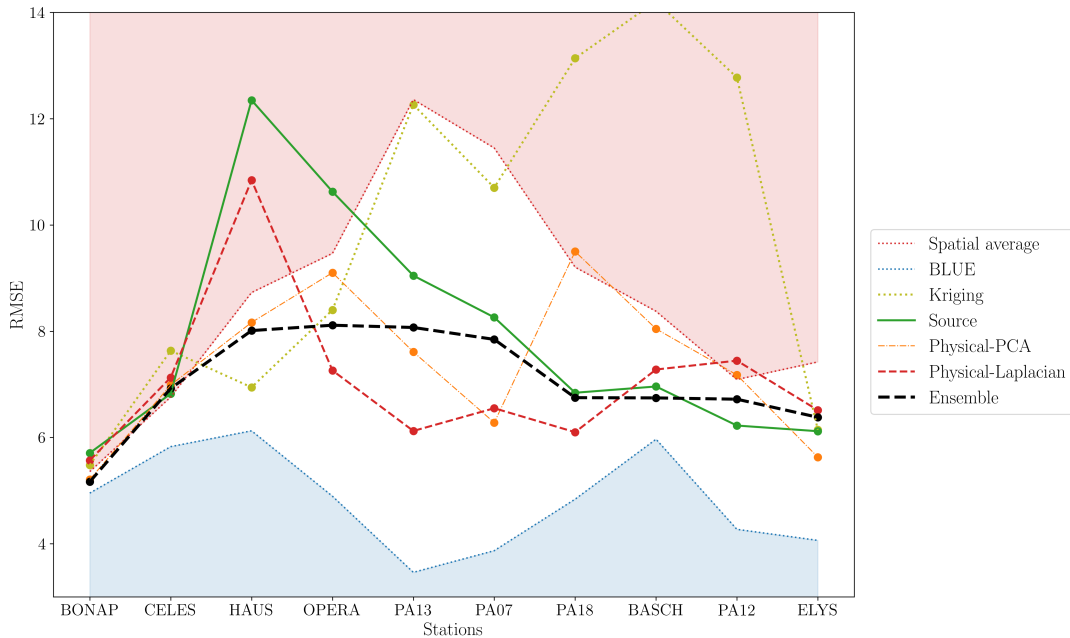


Figure 5: Root mean square error on tested stations for the different proposed methods

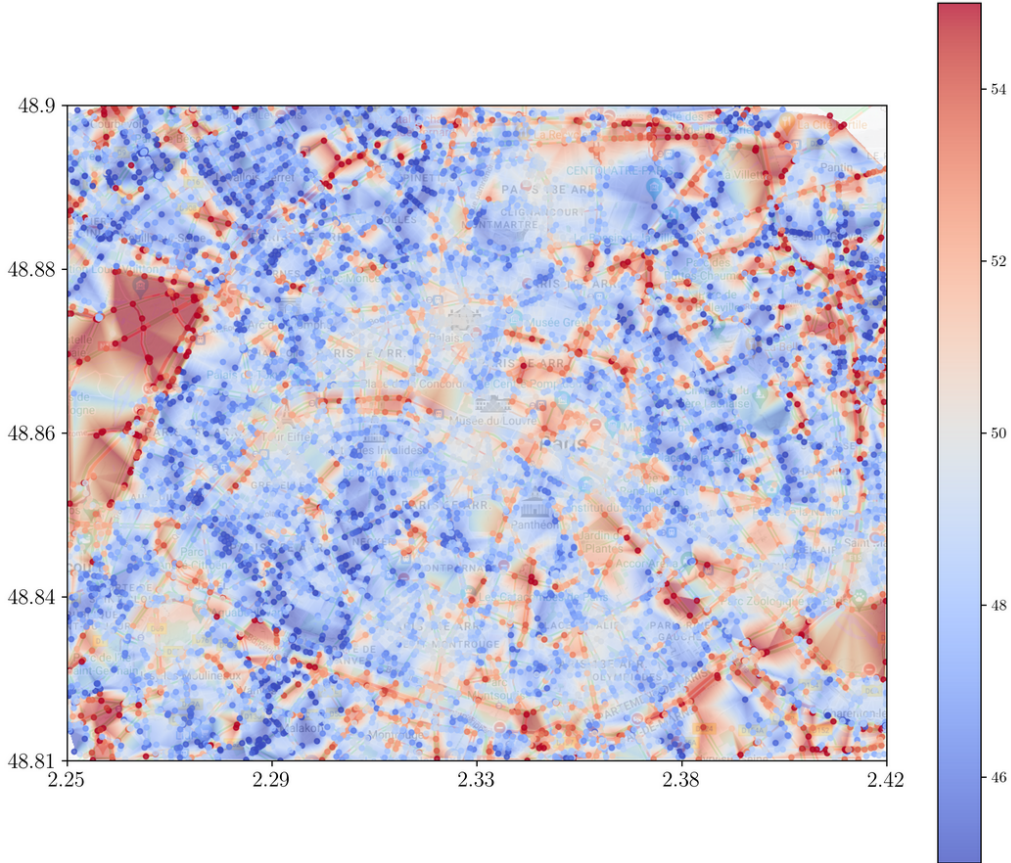


Figure 6: Pollution map for the *ensemble* model at 8am on March 1st, 2023. Based on the predictions on the node, one can linearly extrapolate pollution values even outside of the graph edges. Note that the fine variations in pollutant concentration (between 45 and 55 $\mu\text{g}/\text{m}^3$) seem to trace the main circulation axes.

6 Conclusion and future works

In this work, we have showed that it is possible to leverage pollution sensor data, meteorological information and Google Traffic images to create pollution maps in real time. In particular, we explained how to build statistical, physics-based and ensemble reconstruction strategies by posing the problem of pollution state estimation on metric and quantum graphs. Furthermore, the right combination of these techniques produced systematic improvements over the proposed baselines, namely the *Spatial average* and *Kriging*.

Neither our linear reconstruction strategies nor our nonlinear ones could beat the BLUE benchmark that indicates the accuracy of the best linear estimator. We conjecture that this is due to the limited amount of stations giving us spatial information on the pollution field, and to the indirect nature of traffic data. Regarding the last point, even though the volume of traffic information is large, it still remains of reduced utility. This is due to the fact that we only measure the fluidity of traffic, instead of the actual amount of passing vehicles, which is the relevant variable directly impacting emissions. One can then hope to obtain further improvements by following a similar approach with better suited data.

Another limitation is the unavailability of local pollution averages for the city of Paris. Having access to this kind of data through intensive measurement campaigns lasting a few weeks, but employing hundreds to thousands of sensors, as done in [10] for the city of Barcelona,

would give a much more precise baseline, and allow to learn corrections to the local average instead of the global one.

Finally, in setting the problem on the graph, we did not take into account the vicinity of open spaces like parks or rivers, nor the topology and variations in altitude. This is particularly visible in fig. 6, where the parks of Boulogne and Vincennes are colored in red because of the surrounding highways, and the absence of small internal streets. We leave the inclusion of such relevant features to future studies.

Acknowledgments and disclosure of funding: The authors would like to thank Prof. Albert Cohen, Prof. Joubine Aghili, Dr. Rachida Chakir, Dr. Vivien Mallet, and Dr. Fabien Brocheton for fruitful discussions, and for preliminary work on the extraction of applicable data.

This work was done in the framework of the research project “Models and Measures” funded by the Parisian City Council (Emergences grant program). In addition, Matthieu Dolbeault acknowledges funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - Project number 442047500 through the Collaborative Research Center “Sparsity and Singular Structures” (SFB 1481).

A Metric graphs

Here we recall several notions about graphs that are necessary in our developments. The presentation is based on the book [14], which provides a comprehensive introduction to quantum graphs, and on the paper [18], which develops finite element discretizations of elliptic operators in quantum graphs. We sometimes narrow down the generality of certain notions for the purposes of the present paper.

A combinatorial graph $G = (V, E)$ is a collection of a finite number of vertices V and of edges $E \subset V \times V$ connecting pairs of vertices. We restrict our attention to undirected graphs where no orientation is assigned to the edges, and denote $|V|$ and $|E|$ the number of vertices and edges, respectively.

We will work with connected graphs, where any two vertices $v, w \in V$ are connected by at least one path $(v, v_1), (v_1, v_2), \dots, (v_k, w)$ made by consecutive adjacent edges in E . A connected graph becomes a metric graph if we assign a length $\ell_e > 0$ and a local coordinate $r_e(x)$, for $x \in [0, \ell_e]$, to each edge $e = (v, w) \in E$, in such a way that $r_e(0) = v$ and $r_e(\ell_e) = w$.

In our case, the crossroads V are embedded in \mathbb{R}^2 through their geographical coordinates, and the streets $r_e([0, \ell_e]) \subset \mathbb{R}^2$ are differentiable curves with no loops. However, as done very often, we redefine them as simple straight lines joining the two vertices. Regardless of the choice of the edge curves, the points r in a metric graph G are thus not only its vertices but also all intermediate points on the edges as well, parametrized by the local coordinates r_e :

$$V \subsetneq G = \bigcup_{e \in E} r_e([0, \ell_e]).$$

As the name suggests, any metric graph can be endowed with a natural metric as follows. The distance between two vertices $v, w \in V$ is as usual defined as the length of the shortest path connecting them. This notion of distance between vertices is then extended in a natural way to any two points possibly lying on different edges, by further adding the local coordinates along these edges.

We may now introduce function spaces and linear differential operators on a metric graph G . The space of continuous functions $\mathcal{C}(G)$ contains the functions $u : G \rightarrow \mathbb{R}$ such that $u \circ r_e$ is

continuous on $[0, \ell_e]$ for any edge $e \in E$, which implies in particular the continuity of u along any path in G . The space of square-integrable functions

$$L^2(G) = \bigoplus_{e \in E} L^2(r_e([0, \ell_e]))$$

is a Hilbert space when endowed with the inner product

$$\langle u, v \rangle_{L^2(G)} := \int_G u(r)v(r)dr = \sum_{e \in E} \int_0^{\ell_e} u(r_e(x))v(r_e(x))dx.$$

Finally, the Sobolev space

$$H^1(G) = C(G) \cap \bigoplus_{e \in E} H^1(r_e([0, \ell_e]))$$

is also a Hilbert space, for the norm

$$\|u\|_{H^1(G)}^2 := \int_G u^2 dr + \int_G \left(\frac{du}{dr} \right)^2 dr = \sum_{e \in E} \int_0^{\ell_e} u(r_e(x))^2 dx + \int_G \left(\frac{d(u \circ r_e)}{dx} \right)^2 dx.$$

The restriction to $C(G)$ in the definition of $H^1(G)$ stems from the fact that functions in $H^1(r_e([0, \ell_e]))$ are continuous (because their domain is one dimensional), which automatically implies that functions in $H^1(G)$ must be continuous also at the vertices. In the same vein, one has to impose restrictions on the derivatives of u , such as Neumann-Kirchoff boundary conditions, for functions $u \in H^2(G)$.

References

- [1] Thomas J Santner, Brian J Williams, William I Notz, and Brian J Williams. The design and analysis of computer experiments, volume 1. Springer, 2003.
- [2] Noel Cressie. Statistics for spatial data. John Wiley & Sons, 2015.
- [3] J. S. Hesthaven, G. Rozza, and B. Stamm. Certified reduced basis methods for parametrized partial differential equations. SpringerBriefs in Mathematics, 2015.
- [4] O. Mula. Inverse problems: A deterministic approach using physics-based reduced models. In Model Order Reduction and Applications: Cetraro, Italy 2021, pages 73–124. Springer Nature Switzerland, Cham, 2023.
- [5] Roberto Molinaro, Yunan Yang, Björn Engquist, and Siddhartha Mishra. Neural inverse operators for solving PDE inverse problems. arXiv preprint arXiv:2301.11167, 2023.
- [6] Wolfgang Dahmen, Min Wang, and Zhu Wang. Nonlinear reduced DNN models for state estimation. arXiv preprint arXiv:2110.08951, 2021.
- [7] L. Breiman. Stacked regressions. Machine learning, 24:49–64, 1996.
- [8] M. J. Van der Laan, E. C. Polley, and A. E. Hubbard. Super learner. Statistical applications in genetics and molecular biology, 6(1), 2007.
- [9] E. C. Polley and M. J. Van der Laan. Super learner in prediction. bepress, 2010.

- [10] Alvaro Criado, Jan Mateu Armengol, Hervé Petetin, Daniel Rodríguez-Rey, Jaime Benavides, Marc Guevara, Carlos Pérez García-Pando, Albert Soret, and Oriol Jorba. Data fusion uncertainty-enabled methods to map street-scale hourly NO₂ in Barcelona: a case study with CALIOPE-Urban v1.0. Geoscientific Model Development, 16(8):2193–2213, April 2023.
- [11] Mingfei Niu, Yufang Wang, Shaolong Sun, and Yongwu Li. A novel hybrid decomposition-and-ensemble model based on CEEMD and GWO for short-term PM_{2.5} concentration forecasting. Atmospheric Environment, 134:168–180, June 2016.
- [12] V. Mallet, G. Stoltz, and B. Mauricette. Ozone ensemble forecast with machine learning algorithms. Journal of Geophysical Research: Atmospheres, 114(D5), 2009.
- [13] A. Tilloy, V. Mallet, D. Poulet, C. Pesin, and F. Brocheton. BLUE-based NO₂ data assimilation at urban scale. Journal of Geophysical Research, 118(4):2,031–2,040, 2013.
- [14] G. Berkolaiko and P. Kuchment. Introduction to quantum graphs, volume 186. American Mathematical Soc., 2013.
- [15] Junjie Cai, Jingtian Chen, Haimei Cheng, Shuangfei Zi, Jinchao Xiao, Fan Xia, and Jiyun Zhao. The effects of thermal stratification on airborne transport within the urban roughness sublayer. International Journal of Heat and Mass Transfer, 184:122289, 2022.
- [16] Jianghanyang Li, Xuan Zhang, John Orlando, Geoffrey Tyndall, and Greg Michalski. Quantifying the nitrogen isotope effects during photochemical equilibrium between NO and NO₂: implications for $\delta^{15}\text{N}$ in tropospheric reactive nitrogen. Atmospheric Chemistry and Physics, 20(16):9805–9819, 2020.
- [17] Ralf Kurtenbach, Jörg Kleffmann, Anita Niedojadlo, and Peter Wiesen. Primary NO₂ emissions and their impact on air quality in traffic environments in germany. Environmental Sciences Europe, 24:1–8, 2012.
- [18] M. Arioli and M. Benzi. A finite element method for quantum graphs. IMA Journal of Numerical Analysis, 38(3):1119–1163, 2018.
- [19] A. Cohen and R. DeVore. Approximation of high-dimensional parametric PDEs. Acta Numer., 24:1–159, 2015.
- [20] A. Cohen, R. DeVore, and C. Schwab. Analytic regularity and polynomial approximation of parametric and stochastic elliptic PDE’s. Anal. Appl., 09(01):11–47, 2011.
- [21] J. L. Eftang, A. T. Patera, and E. M. Rönquist. An “hp” certified reduced basis method for parametrized elliptic partial differential equations. SIAM J. Sci. Comput., 32(6):3170–3200, 2010.
- [22] J. S. Hesthaven, G. Rozza, and B. Stamm. Certified reduced basis methods for parametrized partial differential equations. SpringerBriefs in Mathematics, 2015.
- [23] G. Rozza, D. B. P. Huynh, and A. T. Patera. Reduced basis approximation and a posteriori error estimation for affinely parametrized elliptic coercive partial differential equations. Arch. Comput. Methods Eng., 15(3):229–257, 2008.
- [24] Z. Zou, D. Kouri, and W. Aquino. An adaptive local reduced basis method for solving PDEs with uncertain inputs and evaluating risk. Comput. Methods Appl. Mech. Engrg., 345:302–322, 2019.

- [25] Y. Maday, A. T. Patera, J. D. Penn, and M. Yano. A parameterized-background data-weak approach to variational data assimilation: formulation, analysis, and application to acoustics. Internat. J. Numer. Methods Engrg., 102(5):933–965, 2015.
- [26] P. Binev, A. Cohen, W. Dahmen, R. DeVore, G. Petrova, and P. Wojtaszczyk. Data assimilation in reduced modeling. SIAM/ASA J. Uncertain. Quantif., 5(1):1–29, 2017.
- [27] A. Cohen, W. Dahmen, R. DeVore, J. Fadili, O. Mula, and J. Nichols. Optimal reduced model algorithms for data-based state estimation. SIAM J. Numer. Anal., 58(6):3355–3381, 2020.
- [28] A. Cohen, W. Dahmen, O. Mula, and J. Nichols. Nonlinear reduced models for state and parameter estimation. SIAM/ASA J. Uncertain. Quantif., 10(1):227–267, 2022.
- [29] A. Cohen, M. Dolbeault, O. Mula, and A. Somacal. Nonlinear approximation spaces for inverse problems. Anal. Appl., 21(1):217–253, 2023.
- [30] J. K. Hammond, R. Chakir, F. Bourquin, and Y. Maday. PBDW: A non-intrusive reduced basis data assimilation method and its application to an urban dispersion modeling framework. Appl. Math. Model., 76:1–25, 2019.