Bayesian modeling for spatially misaligned health and air pollution data through the INLA-SPDE approach

Michela Cameletti[a,*], Virgilio Gómez Rubio[b], Marta Blangiardo[c]

[a]*Dept. of Management, Economics and Quantitative Methods, University of Bergamo, Italy*
[b]*Dept of Mathematics, Universidad de Castilla-La Mancha, Spain*
[c]*MRC Centre for Environment and Health, Department of Epidemiology and Biostatistics, Imperial College London*

Abstract

In air pollution studies a key issue concerns the change of support: pollutant concentrations are continuous phenomena in space but their measurements are typically available at a finite number of point-referenced monitoring stations or result from numerical models. When linking exposure to health outcomes, the latter are usually available at administrative level, hence on an irregular lattice, providing challenges in terms of data misalignment.

In this paper we tackle the change of support problem for air pollution and health studies through a two-stage Bayesian approach; in the first stage our model estimates the air pollution concentration at the area level and then in the second stage it links the exposure to the health outcome, accounting for the uncertainty on the exposure estimates. We show through an extensive and realistic simulation that our model is able to predict the concentration accurately at the administrative level as well as estimate the association between exposure and health outcome. We use the Integrated Nested Laplace Approximation, coupled with the Stochastic Partial Differential Equation method for model implementation. Finally we apply the proposed model to evaluate the effect of NO_2 concentration on hospital admissions for respiratory diseases in the Piemonte region (Italy). We found that the upscaling method and the approach used to propagate uncertainty from the first to the

*Corresponding author
Email addresses: michela.cameletti@unibg.it (Michela Cameletti), virgilio.gomez@uclm.es (Virgilio Gómez Rubio), m.blangiardo@imperial.ac.uk (Marta Blangiardo)

second stage has an impact on the posterior distribution of the relative risk. Moreover, we found a significant increased risk of 1.6% and 1.8% associated to an increase of 10 $\mu g/m^3$ in NO_2 concentration.

Keywords: Spatial misalignment; Integrated Nested Laplace Approximation (INLA); Stochastic Partial Differential Equations (SPDE); Hierarchical modeling; Uncertainty propagation; Air pollution

1. Introduction

Air pollution is both an environmental and social criticality and it represents the single largest environmental health risk in Europe today (Lim et al., 2012). The recent report of the European Environmental Agency (EEA, 2017a) states that 19% of the urban population in the EU-28 was exposed in 2015 to PM_{10} (particulate matter with an aerodynamic diameter of less than 10 μm) concentrations above the EU daily limit value of 50 $\mu g/m^3$; the same percentage rises to 53% if the WHO stricter threshold (set to 20 $\mu g/m^3$) is considered. The same report states that 9% of the EU-28 urban population lived in areas with concentrations of NO_2 (nitrogen dioxide) exceeding the annual EU limit value of 40 $\mu g/m^3$ in 2015.

In terms of health impact, in 2014 long-term exposure to $PM_{2.5}$ was responsible for about 428,000 premature deaths in Europe (of which around 399,000 were in the EU-28), mainly due to heart and lung diseases (EEA, 2017a), while NO_2 concentration accounted for about 78,000 premature deaths per year (about 75,000 in the EU-28). A large number of epidemiological studies have shown short and long term effects of air pollution on mortality (see for instance Raaschou-Nielsen et al. 2012; Faustini et al. 2014; Atkinson et al. 2016; Halonen et al. 2016; Carugno et al. 2016) or hospital admissions (among the others see for instance Halonen et al. 2016; Carey et al. 2016; Sanyal et al. 2018). Recently some work has appeared suggesting even a link with drug prescriptions for chronic diseases like asthma and COPD, in a primary care perspective (Blangiardo et al., 2016; Lee, 2018). Air pollution has also a considerable economic impact in terms of increased medical costs, reduced productivity and decrease in crop yields: the OECD estimates that these costs will gradually increase to 1% of global worldwide GDP (around USD 2.6 trillion annually) by 2060 (OECD, 2016).

Substantive methodological work has been published to assess the presence of health effects associated with air pollution, including cohort, time

series and small area studies (e.g. see the recent review by Bruno et al., 2016). In this paper we focus on the latter case and consider ecological spatial regression models for aggregated health data consisting of mortality or morbidity counts at the small area level (typically administrative, e.g. electoral wards, district, etc.) together with pollutant measurements available for a set of monitoring stations or grid cell centroids.

In this modeling framework the first statistical challenge regards the spatial misalignment between health and exposure data, with the consequence that pollutant concentration has to be upscaled at the area level, while being measured at a finite number of point-referenced monitoring stations, a procedure known as *change of support* (Gelfand, 2010). The simplest solution consists in averaging the concentration point measurements available for each area, possibly using distance- or population-based weights (e.g. Elliott et al., 2007; Madsen et al., 2008; Young et al., 2009). However this approach is not feasible where the network is sparse, i.e. there are areas without monitoring stations, or when the considered pollutants show a strong spatial heterogeneity. As a solution, it is possible to estimate the area-level concentration by computing weighted averages over grid level concentrations available from deterministic atmospheric dispersion models (see e.g. Bell, 2006; Rushworth et al., 2014; Lee and Sarran, 2015). Alternatively, in a model-based perspective, a spatial statistical model can be built which combines pollutant measurements from stations with the output of numerical dispersion models in a data fusion approach (see for instance Fuentes et al., 2006; Peng and Bell, 2010; Sahu et al., 2010; Berrocal et al., 2010; Pannullo et al., 2015; Moraga et al., 2017). Under this modeling framework, beside accounting for spatial correlation and measurement error, it is easy to include additional covariates like meteorological variables. In this paper we adopt a model-based approach and we deal with the spatial misalignment by computing concentration at the area level through a weighted mean with two different types of weights (linear combination with neighbourhood intersections or simple mean). This approach requires to use the exposure model to obtain predictions for a set of points belonging to a regular grid which covers the region of interest.

The second statistical challenge concerns how to link the exposure with the health outcomes. Commonly a *two-stage* approach is used: the pollutant concentrations estimated at the first stage are then averaged at the level of the irregular lattice where the health data are available; then the posterior mean or median for each area is included as covariate in the second stage (see e.g. Lee and Shaddick, 2010; Huang et al., 2015; Lee et al., 2015; Pannullo

et al., 2016; Liu et al., 2016). The advantage of this approach is mainly computational since the exposure and health models are fitted separately. The crucial issue with a two-stage approach is that it typically treats the area predicted exposures as known and constant, without accounting for the uncertainty in the prediction of the first stage. This may result in overprecise estimates of the risk effect associated to air pollution concentration, or even biased results if the exposure is also assumed to suffer from measurement error.

Some solutions have been recently proposed in the literature and regard the propagation of the uncertainty from the first to the second stage through: *i)* multiple simulation of exposure values (from the pollutant concentration posterior predictive distributions) followed by iterative fits of the health model (Blangiardo et al., 2016; Liu et al., 2016; Lee et al., 2017); *ii)* considering the exposure as a random variable in the health model with an informative prior obtained from the posterior distributions of the exposure model (Warren et al., 2012; Powell and Lee, 2014; Lee et al., 2017; Huang et al., 2017). However no papers have explicitly compared the performance of these different strategies to account for uncertainty from the exposure into the health model.

In this paper we are framed in the same perspective as Liu et al. (2016) and Lee et al. (2017), and build a two-stage model to predict air pollution at a regular grid and to evaluate its health effects at small (administrative) area level. In particular in the first stage we estimate the pollutant concentration via the integration of data from monitoring stations as well as numerical model output and additional covariates, while in the second stage we link the estimated exposure to the health outcome. The novel aspect of the paper consists in the thorough evaluation of the impact of the averaging from grid to small area on the exposure estimates as well as of the uncertainty propagation from stage 1 to stage 2 of the modeling framework on the health outcomes. In order to do so we develop an extensive and realistic simulation study, which we believe will be useful for other researchers working on environmental and health studies at the area level. We then use the proposed framework to study the relationship between NO_2 concentration and hospitalisations for respiratory causes for each communality of the Piemonte region (Italy) for the year 2011. We implement the two-stage model by means of the Integrated Nested Laplace Approximation (INLA; Rue et al., 2009) and Stochastic Partial Differential equations (SPDE; Lindgren et al., 2011) approach, as a computationally effective alternative to the standard approach based on Markov chain

Monte Carlo methods (MCMC). The R code to reproduce the simulation analysis is available at https://github.com/michelacameletti/INLA_COSP.

The rest of the paper is structured as follows: in section 2 we introduce the case study on air pollution and hospital admissions in the Piemonte region in Italy; in section 3 we present the modeling framework, while section 4 briefly describes the INLA-SPDE approach used for the implementation. Section 5 introduces the simulation design and presents its results, then section 6 presents the results of the real data application on NO_2 and hospital admissions in Piemonte, while in section 7 we raise discussion points and concluding remarks.

2. Motivating problem: air pollution and hospitalizations in Piemonte regions, Italy

Piemonte is located in the North-Western part of Italy (see Figure 1, left). Together with Lombardia, Veneto and Emilia Romagna, it is part of the Po Valley, a densely populated and heavily industrialised area located at the footstep of the Alps and characterized by a wide variety of pollution sources mainly related to traffic, domestic heating, farming activities, etc. Its particular geographic position, with the Alps acting like a shelter, leads to frequent occurrence of stagnant meteorological conditions with absence of wind and reduction of pollutant dispersion. For these reasons, Po Valley has been identified as one of the most polluted European regions where pollutant standards, set for human health protections, are regularly exceeded (EEA, 2017a,b).

For this work we consider the annual NO_2 mean concentrations (in $\mu g/m^3$) for 2011, obtained from the 55 monitoring stations depicted with red (for training sites) and green (for validation sites) points in the left plot of Figure 1. NO_2 values range from 16 to 71.88 $\mu g/m^3$ with a median of 35.63 $\mu g/m^3$. In order to make the distribution of NO_2 approximately normal, we use a logarithmic transformation. In addition, in the exposure model the following covariates are considered: precipitation (in mm), mixing height (in m), temperature (in K), windspeed (in m/s) and NO_2 emissions (in g/s), that are obtained from a nested system of deterministic computer-based models implemented by the environmental agency *ARPA Piemonte*. These deterministic models provide data at the monitoring station sites and for all the points of a 4 km $\times$ 4 km grid covering Piemonte region (Cameletti et al.,

2011, 2013). This grid, which has a resolution of $56 \times 72 = 4032$ points (see the blue points in Figure 1, left) will be used for spatial prediction.

To assess the impact of air pollution exposure on human health, we consider hospitalizations data provided by the Ministry of Health. For each patient discharged from either a public or private healthcare facility, we have data about socio-demographic variables (e.g. gender, age) and the hospitalization event (e.g., diagnosis, dates of admission and discharge). In this paper we consider cardio-respiratory hospitalizations occurred in Piemonte region during 2011 and aggregated at the municipality level. The Standardized Morbidity Ratio (adjusted by age and gender) is reported in the right plot of Figure 1 and shows substantial spatial variability, with higher risks around large cities such as Turin and Alessandria, while the more rural central part as well as most of the mountain region is characterized by risks lower than averages.

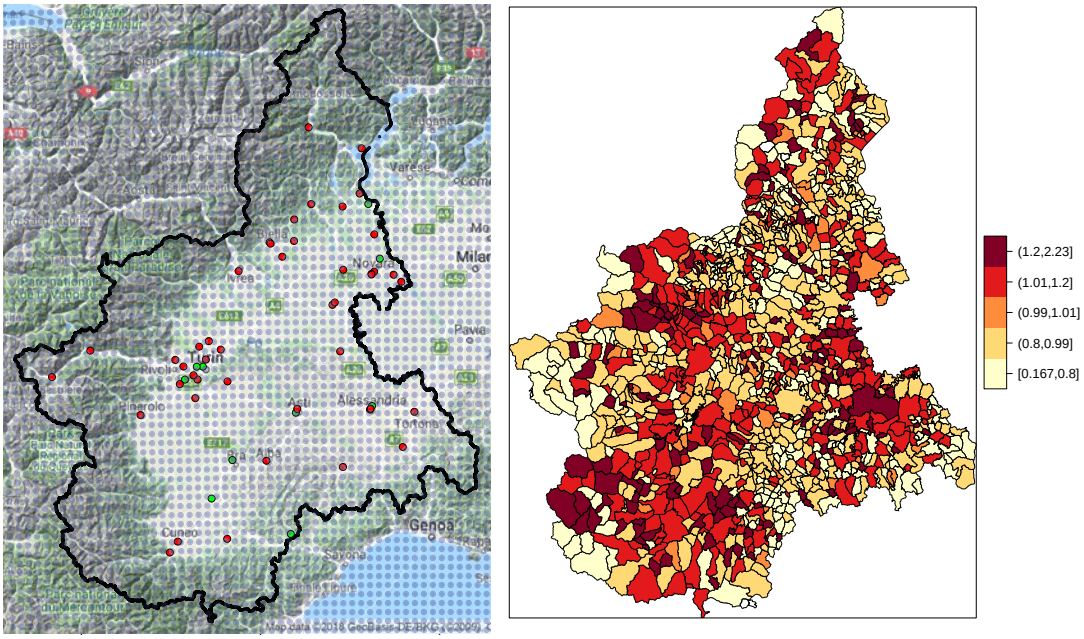

Figure 1: Left: Piemonte map: the red and green dots denote the NO_2 monitoring stations used for model estimation and validation, respectively. The blue dots represent the centroids of the regular grid. Right: Standardized Morbidity Ratio (adjusted by age and gender) for cardio-respiratory hospitalizations in the Piemonte municipalities in 2011.

3. Poisson health model

The standard spatial model for the observed number of health events y_i in the area $\mathcal{A}_i$ ($i = 1, \ldots, n_{\mathcal{A}}$), when there is a relatively low count of disease and/or the area is small, is given by

$$y_i \sim \text{Poisson}(E_i \rho_i)$$

where E_i represents the expected number of events obtained applying standard rates from the whole study region (by age and gender) to the population of each area (Shaddick and Zidek, 2016). The term ρ_i represents the area specific relative risk and a linear predictor is defined on its logarithmic transformation as follows:

$$\log(\rho_i) = \gamma_0 + x_i \gamma_1 + \boldsymbol{z}_i' \boldsymbol{\gamma}_2 + \phi_i. \tag{1}$$

This linear predictor includes an intercept γ_0 (representing the average rate in the entire study region), the area air pollution concentration x_i, random effects ϕ_i and a vector of measured confounders $\boldsymbol{z}_i'$. In the application presented in Section 6 we include the social vulnerability index provided by the Italian National Institute of Statistics (ISTAT) for 2011, which measures the deprivation level of individuals within municipalities. To evaluate the health risk associated to air pollution, the parameter of interest is γ_1 or the corresponding relative risk given by $\exp(\gamma_1 \delta)$, which represents the change in the risk of experiencing the considered health outcome when air pollution concentrations increases by δ $\mu g/m^3$.

The random effects ϕ_i capture any overdispersion and potential residual spatial correlation in the health data after the covariate effects have been accounted for. An additive specification can be adopted such that $\phi_i = u_i + v_i$, and several structures can be assumed on these two terms (see Lee, 2011 for a review). In the application of Section 6 we use the specification by Besag et al. (1991), which places an exchangeable random effect on v_i and a conditional autoregressive structure on u_i so that

$$\begin{aligned} v_i &\sim \text{Normal}(0, \sigma_v^2) \\ u_i \mid u_{-i} &\sim \text{Normal}\left(\frac{\sum_{j \in D_i} u_j}{|D_i|}, \frac{\sigma_u^2}{|D_i|} \right), \end{aligned}$$

where D_i represents the set of areas sharing borders with the i-th area and $|D_i|$ its cardinality. This assumes that only the areas close to each other can

influence one another and provides some local smoothing to the estimates of the relative risks. In addition some global smoothing is provided by the spatially unstructured random effects v_i.

3.1. Exposure estimation through upscaling

The term x_i in Eq. (1) represents the exposure level for area $\mathcal{A}_i$; it can not be measured directly as air pollutant concentrations are available only for a finite number of spatial points with coordinates $\boldsymbol{s}_1, \ldots, \boldsymbol{s}_n$. The set of point-referenced concentration measurements is denoted by $(x(\boldsymbol{s}_1), \ldots, x(\boldsymbol{s}_n))$ and is a realization of the latent stochastic process $x(\boldsymbol{s})$ representing the true air pollution field, which is continuous in space. The average exposure level for area $\mathcal{A}_i$ would be given by

$$x_i = \int_{\boldsymbol{s} \in \mathcal{A}_i} x(\boldsymbol{s})\, p(\boldsymbol{s})\, d\boldsymbol{s} \tag{2}$$

where $p(\boldsymbol{s})$ is a weight for a generic spatial point $\boldsymbol{s} \in \mathcal{A}_i$ such that $\int_{\boldsymbol{s} \in \mathcal{A}_i} p(\boldsymbol{s})\, d\boldsymbol{s} = 1$ (Gelfand, 2010).

This stochastic integral can be estimated using the set of measurements from the n monitoring stations by simply averaging the concentration data from the stations falling within each area $\mathcal{A}_i$. However, given that monitoring networks are typically sparse, some areas could end up with no monitoring stations; a possible solution would consist in estimating the concentration only within a specific distance from monitoring stations and then evaluating the health effects only on the population within the same areas (as in Zhu et al. 2003). Alternatively, in order to cover all the spatial domain, the area exposure level x_i can be computed through Monte Carlo integration using a set of additional points, denoted by $\boldsymbol{s}^\star$, which are the centroids of a regular grid covering the region of interest (see e.g. Lee and Shaddick, 2010; Lee and Sahu, 2016). With this approach the exposure value for area $\mathcal{A}_i$ is estimated through the following weighted mean:

$$x_i = \sum_{j=1}^{N_i} x(\boldsymbol{s}^\star_{ij})\, p(\boldsymbol{s}^\star_{ij}), \tag{3}$$

where $x(\boldsymbol{s}^\star_{ij})$ is the pollutant concentration value for the generic location $\boldsymbol{s}^\star_{ij}$, which is one of N_i regular grid centroids inside area $\mathcal{A}_i$. The corresponding weights are normalized so that $\sum_{j=1}^{N_i} p(\boldsymbol{s}^\star_{ij}) = 1$. The exposure values $x(\boldsymbol{s}^\star_{ij})$

can: *i)* be provided by an air pollution (deterministic) dispersion model, such as the Community Multiscale Air Quality Modeling System (CMAQ), used by the US Environmental Protection Agency, or the Atmospheric Dispersion Modeling System (ADMS), particularly useful for urban areas; *ii)* be the output of simple spatial interpolation (e.g. inverse distance weighting, kriging) using the monitoring network data; *iii)* be derived by spatial prediction from an *exposure model* which can fuse different sets of data, account for the measurement error and include explicitly a continuous spatial process (see e.g. Sahu, 2011; Lee et al., 2017) as described in Section 3.2; in the Bayesian framework, this means that an exposure posterior predictive distribution is available for each grid point $s^{\star}_{ij}$ and it can be used to derive the area level exposure posterior distribution. In particular, the prediction of the exposure at the area level is performed in two steps: firstly the exposure posterior distribution is obtained for a set of points $\boldsymbol{s}^{\star}_{ij}$ belonging to a regular grid. Secondly, for each area $\mathcal{A}_i$, the exposure average is computed using two methods:

1. Method 1 (*linear combination with neighbourhood intersections*): the area exposure is computed using Eq. (3), with $x(\boldsymbol{s}^{\star}_{ij})$ being the exposure estimates available at the centroid $\boldsymbol{s}^{\star}_{ij}$ of the N_i cells which have an intersection with the considered area $\mathcal{A}_i$ ($j = 1, \ldots, N_i$). Note that we assume that the generic weight $p(\boldsymbol{s}^{\star}_{ij})$ is given by the proportion of the j-th grid cell overlapping with area $\mathcal{A}_i$. For example, Figure 2 (left) represents a generic area of the considered region which intersects 11 cells of the regular grid ($N_i = 11$) and with corresponding weights ranging from 0.001 to 0.236.
2. Method 2 (*simple mean*): the area exposure is computed using Eq. (3), but considering $x(\boldsymbol{s}^{\star}_{ij})$ as the exposure estimates available at the grid cell centroids $\boldsymbol{s}^{\star}_{ij}$ that lie in the considered area. If no grid centroids are located inside the area, then the closest grid point is used. In this case we consider a system of equal weights for all the considered grid points used in the linear combination. For example, for the generic area shown in Figure 2 (right) 3 grid points lie inside the area and their weights are all equal to 1/3.

It would also be possible to estimate the area exposure x_i by averaging grid predictions with weights $p(\boldsymbol{s}^{\star}_{ij})$ proportional to the population at-risk (e.g. Wakefield and Shaddick, 2006). This approach requires to have high-resolution information about the population size (for example the LandScan™ project provides global population estimates at 1km spatial resolution) and

to perform some geoprocessing to align the spatial datasets (e.g. Shaddick et al., 2018).

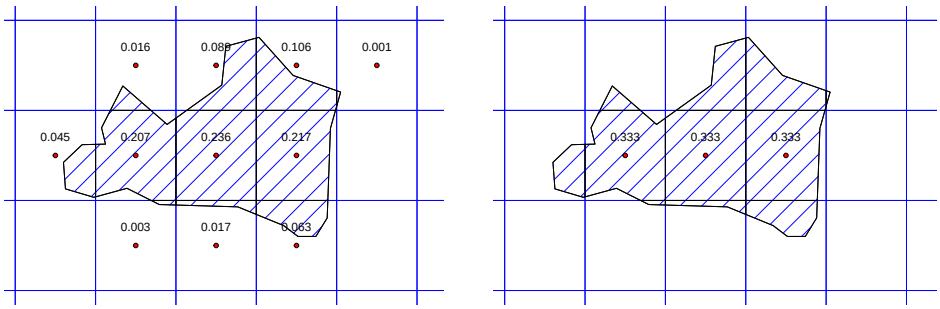

Figure 2: Method1: intersections between a generic area and the regular grid and corresponding weights (left). Method2: grid points inside a generic area and corresponding weights (right).

3.2. Gaussian exposure model

Let $z(\boldsymbol{s}_j)$ $(j = 1, \ldots, n)$ be the set of air pollution data measured by n monitoring stations. The pollutant concentration is assumed to be centred on the true concentration $x(\boldsymbol{s}_j)$ and depends on the measurement error $e(\boldsymbol{s}_j)$:

$$z(\boldsymbol{s}_j) = x(\boldsymbol{s}_j) + e(\boldsymbol{s}_j), \tag{4}$$

where $e(\boldsymbol{s}_j) \sim \text{Normal}(0, \sigma_e^2)$ independently for each location. The true concentration is defined by the following linear predictor

$$x(\boldsymbol{s}_j) = b_0 + \boldsymbol{v}(\boldsymbol{s}_j)'\boldsymbol{b} + \omega(\boldsymbol{s}_j) \tag{5}$$

which includes an intercept b_0 (i.e. the average level of pollution for the considered area), a set of site specific covariates $\boldsymbol{v}(\boldsymbol{s}_j)$ (e.g. meteorological and geographical variables) with $\boldsymbol{b}$ vector of coefficients (assumed to be site-invariant) and a latent process $\omega(\boldsymbol{s}_j)$ representing the residual spatial field. The n-dimensional process $\boldsymbol{\omega} = (\omega(\boldsymbol{s}_1), \ldots, \omega(\boldsymbol{s}_n))$ is assumed to be Normally distributed with zero mean vector and spatially structured covariance matrix which is defined by the following Matérn covariance function

$$Cov(\omega(\boldsymbol{s}_j), \omega(\boldsymbol{s}_{j'})) = \frac{\sigma_\omega^2}{\Gamma(\lambda)2^{\lambda-1}} \left(\kappa||s_j - s_{j'}||\right)^\lambda K_\lambda \left(\kappa||s_j - s_{j'}||\right), \tag{6}$$

where $||\boldsymbol{s}_j - \boldsymbol{s}_{j'}|| \in \mathbb{R}$ is the Euclidean spatial distance, σ^2_ω is the spatial variance and κ is the scaling parameter. The term $K_\lambda(\cdot)$ denotes the modified Bessel function of second kind and order $\lambda > 0$. The parameter λ, which is usually kept fixed, measures the degree of smoothness of the process and its integer value determines the mean square differentiability of the process. Instead, $\kappa > 0$ is a scaling parameter related to the range r, i.e. distance at which the spatial correlation is close to 0.1 for each $\lambda \geq 1/2$; Lindgren et al. (2011) proposed an empirically derived definition for the spatial range r, given by $r = \sqrt{8\lambda}/\kappa$. Many other exposure models are available in the literature including also extensions to the spatio-temporal case (see e.g. Sahu, 2011; Cameletti et al., 2011, 2013; Fassò and Finazzi, 2011; Pirani et al., 2014; Pannullo et al., 2015).

In a fully Bayesian approach, we denote by $\boldsymbol{\theta} = \{b_0, \boldsymbol{b}, \boldsymbol{\omega}\}$ the latent Gaussian field and by $\boldsymbol{\psi} = \{\sigma^2_e, \sigma^2_\omega, \kappa\}$ the vector of hyperparameters. This identifies a three-level hierarchical model with the first stage given by observed data distribution $\pi(\boldsymbol{z} \mid \boldsymbol{\theta}, \boldsymbol{\psi})$ with $\boldsymbol{z} = (z(\boldsymbol{s}_1), \ldots, z(\boldsymbol{s}_n))$, the second stage specified by the latent field distribution $\pi(\boldsymbol{\theta} \mid \boldsymbol{\psi})$ and the last level devoted to the hyperparameter prior distribution $\pi(\boldsymbol{\psi})$. Within this modeling framework, the exposure distribution for a new spatial point $\boldsymbol{s}^\star_{ij}$ not included in the set of monitoring stations, is simply given by substituting $\boldsymbol{s}_j$ with $\boldsymbol{s}^\star_{ij}$ in Eq. (4) and (5). The corresponding posterior predictive distribution is then denoted by $\pi(x(\boldsymbol{s}^\star_{ij}) \mid \boldsymbol{z})$ and is given by

$$\begin{aligned} \pi(x(\boldsymbol{s}^\star_{ij}) \mid \boldsymbol{z}) &= \int\int \pi\left(x(\boldsymbol{s}^\star_{ij}), \boldsymbol{\theta}, \boldsymbol{\psi} \mid \boldsymbol{z}\right) \, d\boldsymbol{\theta} d\boldsymbol{\psi} \\ &= \int\int \pi\left(x(\boldsymbol{s}^\star_{ij}) \mid \boldsymbol{\theta}, \boldsymbol{\psi}, \boldsymbol{z}\right) \pi\left(\boldsymbol{\theta} \mid \boldsymbol{\psi}, \boldsymbol{z}\right) \pi\left(\boldsymbol{\psi} \mid \boldsymbol{z}\right) \, d\boldsymbol{\theta} d\boldsymbol{\psi}. \end{aligned} \tag{7}$$

When performing Bayesian inference through Markov chain Monte Carlo methods, samples from the posterior predictive distribution (7) are drawn by composition (Sahu, 2011). In this paper instead we adopt the INLA approach for jointly estimating the parameters and performing spatial prediction both at the grid point and area level.

3.3. Linking the exposure and health model

The easiest and most commonly used method for estimating the adverse effect of air pollution on human health is through a *plug-in* (PI) approach: first the exposure model is estimated (see Section 3.2) and the pollutant

concentration is upscaled at the area level as described in Section 3.1; from this first stage a summary statistic is computed for each area (e.g. exposure posterior mean or median). The second stage consists in including such value as the term x_i in the linear predictor of the health model (see Eq. (1)), which is then fitted separately from the first stage. While being computationally advantageous, this approach does not consider the uncertainty intrinsic in the prediction of the pollutant concentration as only the summary statistics is plugged in. Thus, the resulting risk estimate tends to be unnaturally precise and might even be biased if the concentration suffers from measurement error.

Here we implement two ways of propagating the uncertainty from the exposure into the health model. The first, which we call *feed-forward* (FF) approach, consists in sampling J samples (e.g. $J = 100$) from the joint posterior predictive distribution of the pollutant concentration at the area level and to fit the health model for each one of these simulated exposure values. The posterior distribution of the risk estimate γ_1 (see Eq.(1)) will then be obtained by combining all the results across the J runs. This approach has been adopted also by Blangiardo et al. (2016), Liu et al. (2016) and Lee et al. (2017) and it represents a relatively computationally cheap solution for taking into account the variability of the exposure estimates.

Alternatively, we consider a *prior-exposure* (PE) approach, which specifies an informative prior distribution for the exposure area level x_i in Eq.(1) as follows

$$x_i^A \sim \text{Normal}\left(\mu_i, \sigma_i^2\right) \tag{8}$$

where μ_i and σ_i^2 are given by the posterior means and variances from the area level exposure posterior predictive distributions (see Section 3.1). This corresponds to a multivariate Normal distribution for $x^A = (x_1^A, \ldots, x_{n_A}^A)$ with spatially structured mean vector and diagonal covariance matrix with values given by σ_i^2.

Note that the product of two Gaussian distributions for the term $\gamma_1 x_i^A$ follows by assuming a vague Normal prior distribution also for γ_1.

4. Implementation: the INLA-SPDE approach

We performs Bayesian inference using the integrated nested Laplace approximations (INLA) (Rue et al., 2009; Blangiardo and Cameletti, 2015; Rue et al., 2017), a computationally efficient alternative to MCMC methods for latent gaussian models which can be implemented through the `R-INLA` library (see `http://www.r-inla.org`).

For the first stage of our modeling framework, we couple INLA with SPDE approach proposed by Lindgren et al. (2011) required when Bayesian inference is needed on a spatial process defined over a continuous domain. The SPDE method represents a Gaussian field with Matérn spatial covariance function (see Eq.(6)) as a discrete indexed Gaussian Markov random field (GMRF), which is characterized by a sparse precision matrix and enjoys computational benefits in terms of fast inference. This representation is based on a finite combination of piecewise linear functions defined over a triangulation (or mesh) of the domain of interest and with basis weights defined by a GMRF with sparse precision matrix explicitly depending on the Matérn parameters (Lindgren and Rue, 2015). Spatial prediction in a given location belonging to the considered spatial domain is straightforward since SPDE provides the approximation of the entire spatial process; it is just a matter of including in the INLA model the locations where predictions are required as missing values observations (Lindgren and Rue, 2015).

Note that the PE modeling described in Section 3.3 cannot be run in INLA as it involves the product of two Gaussian distributed parameters (γ and x_i^A) which breaks normality of the latent field. However, by conditioning on γ_1, it is possible to rewrite the product $x_i\gamma_1$ as follows:

$$x_i\gamma_1 = (\mu_i + z_i\sigma_i)\,\gamma_1 = \mu_i\gamma_1 + \gamma_1\sigma_i z_i$$

where $z_i \sim N(0,1)$. This model conditioned on γ_1 can be estimated using INLA: in particular, in the `R-INLA` setting the term $\mu_i\gamma_1$ must be considered as an offset, while the term $\gamma_1\sigma_i$ is the weight of the i.i.d. random effect given by z_i (see Gómez-Rubio and Rue, 2018). Thus, it is possible to obtain the posterior conditional marginals of all the remaining parameters in $\boldsymbol{\theta}$ (including the parameters from the exposure and health model), i.e. $\pi(\boldsymbol{\theta} \mid \gamma_1, \boldsymbol{y})$, and the conditional likelihood $\pi(\boldsymbol{y} \mid \gamma_1)$. To draw values for γ_1 the Metropolis-Hastings (MH) algorithm could be used: after a suitable number of iterations (say L), the MH algorithm will produce samples from $\pi(\gamma_1 \mid \boldsymbol{y})$ denoted by $\{\gamma_1^{(j)}\}_{j=1}^{L}$. Finally, it is possible to get the posterior marginals of all the remaining parameters in $\boldsymbol{\theta}$ by combining all the conditional marginals as follows (see Gómez-Rubio and Rue, 2018):

$$\pi(\theta_i \mid \boldsymbol{y}) = \int \pi(\theta_i \mid \boldsymbol{y}, \gamma_1)\pi(\gamma_1 \mid \boldsymbol{y})\, d\gamma_1 \simeq \frac{1}{L}\sum_{j=1}^{L} \pi(\theta_i \mid \boldsymbol{y}, \gamma_1^{(j)}).$$

4.1. Priors

In `R-INLA` the smoothness parameter λ of the Matérn covariance function in Eq.(6), which is usually kept fixed to ensure model identifiability, is by default equal to 1. The SPDE parameters are represented as $\log(\tau) = \theta_1$ (τ is related to the variance through the relationship $\sigma^2_\omega = 1/(4\pi\kappa^2\tau^2)$) and $\log(\kappa) = \theta_2$, with θ_1 and θ_2 being given independent Normal(0,1) prior distributions (for more details see Blangiardo and Cameletti, 2015). Moreover, weakly informative Normal priors centered on 0 and with a small precision equal to 0.01 are specified for the fixed effects parameters b_0, $\boldsymbol{b}$ and γ_0 and γ_1. Finally all the log precisions are assigned inverse Gamma distributions with parameters equal to 1 and 0.00005.

5. Simulation study

In this section we describe the simulation study which has a twofold aim: firstly, it evaluates the goodness of exposure predictions at the area level obtained through the change of support by using Method 1 and Method 2 (see Section 3.1); secondly it assesses the effect of different ways for incorporating exposure in the health model on the relative risk estimate (i.e. the PI, the FF and the PE approach described in Section 3.3). For the simulation study we use the Belo Horizonte region shapefile, available through the `spdep R` package (Bivand and Piras, 2015); this has a smaller number of areas compared to our Piemonte case study ($n_{\mathcal{A}} = 98$ vs $n_{\mathcal{A}} = 1206$) hence it is more computationally manageable.

5.1. Simulation of the exposure field

In order to create a continuous spatial field, we simulate exposure at a large number m of locations ($m = 4009$) which are aligned in space and cover completely the considered region (see for example left plot in Figure 3). The model used for simulating exposure is based on Eq. (4) and (5) for $j = 1, \ldots, m$. In particular, we assume to have just one covariate v, simulated from a Normal(0,1) distribution, with coefficient $b = 2$ and an intercept equal to $b_0 = 10$. Regarding the spatial parameters (see Eq. (6) and recall that in `R-INLA` $\lambda = 1$ and $\kappa = \sqrt{8}/r$), we set the spatial variance σ^2_ω equal to 0.5 and consider a range r given by 1.63 degrees, corresponding to the 40% of the maximum distance.

Finally, the *true* exposure at the area level is denoted by $\tilde{x}_i$ ($i = 1, \ldots, n_{\mathcal{A}}$) and computed for each area $\mathcal{A}_i$ by averaging the exposure values of the sites located inside the area (see right plot in Figure 3):

$$\tilde{x}_i = \frac{\sum_{\boldsymbol{s}_j \in \mathcal{A}_i} x(\boldsymbol{s}_j)}{m_i}, \tag{9}$$

where m_i denotes the number of sites inside area $\mathcal{A}_i$ ($\sum_i^{n_{\mathcal{A}}} m_i = m$). Note that the cardinality ranges from 7 to 174 points with a median of 30 points per area.

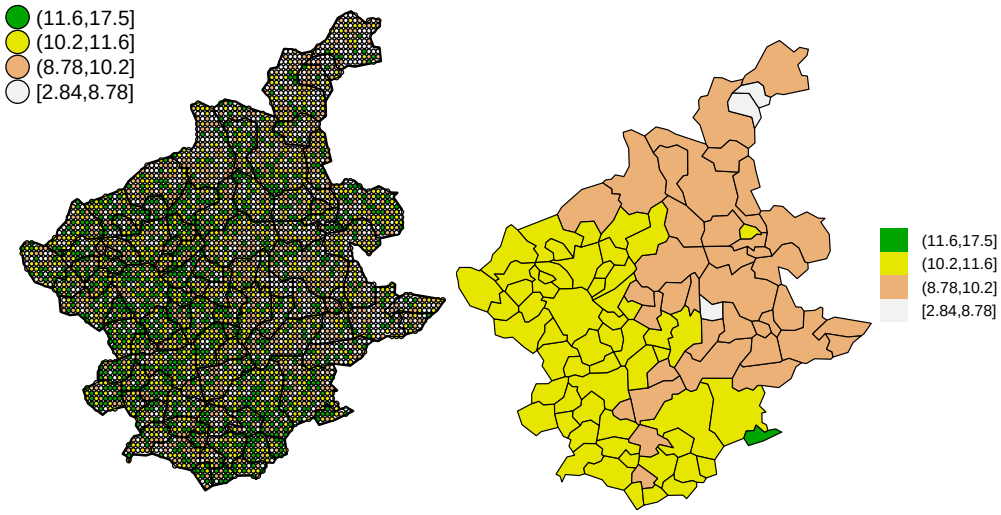

Figure 3: Example of true exposure values simulated at 4009 points inside the Belo Horizonte region (left) and corresponding area averages computed for the 98 areas (right).

5.1.1. Monitoring station sampling

From the set of m spatial locations used for simulating the true exposure field (see the previous Section 5.1), we randomly select n sites which correspond to the monitoring stations where exposure concentration is measured. We consider three cases with the number of stations in each area n_i ($\sum_{i=1}^{n_{\mathcal{A}}} n_i = n$) being 2%, 10% or 30% of the total number of available

points m_i in each area. This leads to a total number of monitoring sites equal to 80, 403 and 1200, respectively. The distribution of points across areas is reported in the left plot of Figure 4: note that the median number of monitoring stations for each area is equal to 1, 3 and 9 for the three cases, respectively.

We assume that the exposure is measured with an error, as specified in Eq. (4). For this reason we add a term to the monitoring station exposure $x(\boldsymbol{s}_j)$, simulated independently from a Normal distribution with mean zero and variance equal to $\sigma_e^2 = 0.1$. This set of data will be used to estimate the exposure field and to predict exposure first at the grid level and then at the area level, as described in the next section.

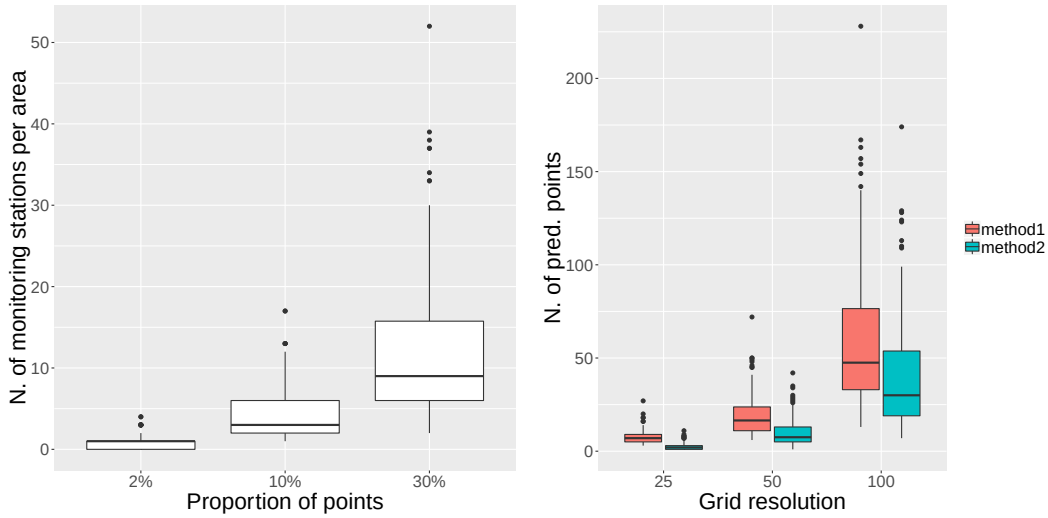

Figure 4: Left: distribution of the number of monitoring stations across areas for the three considered cases (number of monitoring stations equal to 2%, 10% or 30% of the total number of available grid points inside each area, respectively). Right: distribution of the number of prediction points across areas according to the regular grid resolution (25^2, 50^2 and 100^2 points) and the upscaling method (Method 1 and Method 2).

5.1.2. Regular grid for prediction

Spatial prediction of exposure is performed considering a regular square grid with n_g cells covering the entire region and extending also slightly outside (see for example Figure 5 for the case with $n_g = 25^2 = 625$ points). This regular grid is employed for predicting exposure at the grid cell centroids $\boldsymbol{s}_{ij}^{\star}$

using the exposure model described in Section 3.2. As in Section 5.1, the values of the covariate v for the grid cell centroids are drawn independently from a Normal$(0, 1)$ distribution. For the simulation study we consider three grid resolutions with n_g equal to 625, $50^2 = 2500$ and $100^2 = 10000$ points, respectively. Note that the latter corresponds exactly to the grid used for simulating the true exposure field with $m = 4009$ points inside the region (see Section 5.1).

The grid resolution is strictly related to the number of prediction points used for computing the exposure level at the area level using Method 1 and Method 2 described in Section 3.1. The right plot of Figure 4 displays the distribution of the number of prediction points across area according to the regular grid resolution and the upscaling method. It can be observed that, as expected, Method 1 - that considers the intersections between prediction grid and area - employs a higher number of prediction points than Method 2: the median number of points, according to the three grid resolutions, is equal to 7, 16.5 and 47.5 for Method 1 and 2, 7.5 and 30 for Method 2. The total number of prediction points in the whole region is equal to 774, 1957 and 5831 for Method 1 and 257, 1001 and 4009 for Method 2. We expect to be able to predict exposure more accurately at the area level by using a higher number of prediction points, especially for small areas which do not contain any grid square centroids when the grid is coarse.

5.2. Simulation of the health data

Given the true exposure $\tilde{x}_i$ $(i = 1, \ldots, n_A)$ at the area level, it is possible to simulate the health count data y_i using the Poisson model introduced in Section 3. In particular, for the linear predictor of Eq. (3.1) we set $x_i = \tilde{x}_i$ and $\log(\gamma_1)$ equal to $\log(1.05)$, which would be realistic for the impact of air pollution on hospital admissions or mortality in Europe (see for instance Lee and Sarran, 2015; Moore et al., 2016). The corresponding value for the intercept γ_0 is fixed equal to -0.4, a value which guarantees a reasonable spread of the Poisson simulated data. We include also a spatially unstructured random effect $v_i \sim \text{Normal}(0, \sigma_\phi^2)$ with $\sigma_\phi^2 = 0.05$ but for the sake of simplicity omit the spatially structured random effect u_i (see Section 3). The choice of the value for σ_ϕ^2 is done in order to avoid too much variability in the random effect given that the relative risk is small. Finally, we assume that the expected number of cases E_i is fixed and equal to 100 for all the areas (Lee and Sarran, 2015; Wang et al., 2019).

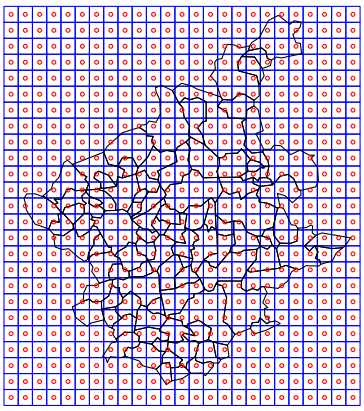

Figure 5: Regular prediction grid with $n_g = 625$.

5.2.1. Simulation scenarios and performance assessment

Combining together the proportions of monitoring stations (n_i equal to 2, 10% and 30% of the sites m_i in each area) with the grid resolution n_g (25^2, 50^2 and 100^2 grid centroids) we obtain 9 scenarios which are summarized in Table 1.

	% of monitoring stations		
n_g	2%	10%	30%
25^2	(1)	(2)	(3)
50^2	(4)	(5)	(6)
100^2	(7)	(8)	(9)

Table 1: Simulation scenarios considering different values for the resolution of the regular grid (n_g) and the percentage of sampled points as monitoring stations.

For the h-th scenario ($h = 1, \ldots, 9$) we run 500 simulations which differ for the true simulated exposure (at the point and area level) and the Poisson data. The monitoring station coordinates are instead fixed across simulations to avoid that their locations influence the estimation and prediction results. Within a specific scenario, the k-th simulation run ($k = 1, \ldots, 500$) is structured as follows:

1. **Simulation**: (i) the exposure field is simulated using $m = 4009$ points and then the true exposure at the area level is computed (see Section 5.1); (ii) the n_i monitoring stations are assigned exposure values equal to the true exposure plus the measurement error realisation, as described in Section 5.1.1; (iii) the health data y_i ($i = 1, \ldots, n_{\mathcal{A}}$) are simulated from the Poisson model using the true area exposure (see Section 5.2).
2. **Estimation and prediction**: (i) Using the INLA-SPDE approach the spatial parameters of the exposure model (b_0, b, σ_e^2, σ_ω^2, r) are estimated and the exposure field is predicted for the grid prediction points. At the same time, the exposure is estimated at the area level using Method 1 and Method 2 for upscaling; (ii) using the three methods for linking the exposure with the health model (PI, FF and PE), the Poisson model with parameters γ_0, γ_1, σ_ϕ^2 is estimated.

Let $\tilde{\theta}$ denote the true value for the generic parameter of interest θ. Given a scenario h and a simulation run k, for each parameter we simulate 100 values, denoted by $\{\hat{\theta}_{hkl}\}$ ($l = 1, \ldots, 100$) from the corresponding posterior distribution. Then for each scenario, simulation and parameter we compute the bias and root mean square error (RMSE) as follows:

$$\text{bias}(\theta) = \frac{1}{100}\sum_{l=1}^{100}\left(\hat{\theta}_{hkl} - \tilde{\theta}\right) \tag{10}$$

$$\text{RMSE}(\theta) = \sqrt{\frac{1}{100}\sum_{l=1}^{100}\left(\hat{\theta}_{hkl} - \tilde{\theta}\right)^2}\ . \tag{11}$$

The same performance indexes are used in order to evaluate the goodness of fit of the exposure predictions for each area $\mathcal{A}_i$, by comparing the true exposure area value $\tilde{x}_{hki}$ and the corresponding estimates $\hat{x}_{hkli}$ which is the l-th value drawn from the exposure posterior predictive distribution of area $\mathcal{A}_i$.

5.2.2. Simulation results

The goodness of fit of the area exposure predictions depends strongly on the number of prediction grid points: as shown in the left and middle plots of Figure 6 the performance indexes (bias and RMSE) are worse for scenarios

1-3 (with 25^2 prediction points), improve for scenarios 4-6 (with 50^2 prediction points) and reach the best values for scenarios 7-9 (with 100^2 prediction points). The latter was expected as the number of prediction points for scenarios 7-9 coincides with the number of points used for simulating the true area exposure ($m = 4009$). Regarding the two upscaling methods (Method 1 and Method 2) it is important to note that Method 2 for scenarios 7-9 represents the benchmark as it is exactly the same method used for simulating the true area exposure (computed as average of the exposure values observed in the prediction points inside each area, with a grid resolution of 100^2 prediction points). The corresponding bias, RMSE and correlation reported in the plots of Figure 6 are not exactly equal to 0 and 1 as expected only due to sampling variability. At the same time, it seems that the performance of Method 1 for scenarios 7-9 is quite similar to the benchmark (Method 2). This holds especially for the bias and RMSE, even if for Method 1 we observe a higher variability of the results. Moreover, for scenarios 1-3 (low resolution prediction grid) Method 1 seems to have larger median biases but lower median RMSEs and higher correlation values. For the remaining scenarios 3-5, the two methods behave very similarly and the indexes' medians basically coincide. Finally, it is worth to note that from the computational point of view Method 1 and Method 2 require the same time to run.

The effect of the number of monitoring stations can be assessed by evaluating differences in the indexes' distribution within groups of scenario 1-3, 4-6 and 7-9: while it seems uninfluential for the prediction bias, increasing the number of monitoring stations helps improve the area predictions in terms of RMSE and correlation, especially within the scenarios 7-9.

The number of monitoring stations has an effect also on the bias and RMSE of the spatial model parameters (b_0, b_1, σ_e^2, σ_ω^2 and r) since the accuracy of the posterior distributions increases when more locations are used for estimation. As expected, there is no effect of the prediction size grid and the parameter estimates for scenario 1-4-7, 2-5-8, 3-6-9 coincide. The results reported in Table A.1 show that the RMSE is lower for the scenarios 3, 6 and 9, which are the ones with the highest number of stations. This is not true only for the spatial variance σ_ω^2 which shows the lowest RMSE for the intermediate scenarios (2, 5 and 8). Regarding the bias, we observe that, on average, we overestimate the spatial variance σ_ω^2 and the range r, with the lowest values obtained in the intermediate scenarios (2, 5 and 8) with 403 monitoring stations. It is worth noting that in general the spatial variance and the range are the most difficult parameters to be estimated and more

informative prior could be adopted to help the inferential procedure (Bakar and Sahu, 2015). For the remaining parameters the values of the bias are quite small, especially for b_1 and σ_e^2. Overall the mean bias (RMSE), averaged across all the scenarios, is -0.022 (0.9) for b_0, -0.001 (0.033) for b_1, -0.003 (0.019) for σ_e^2, 0.179 (0.676) for σ_ξ^2 and 0.19 (1.115) for r.

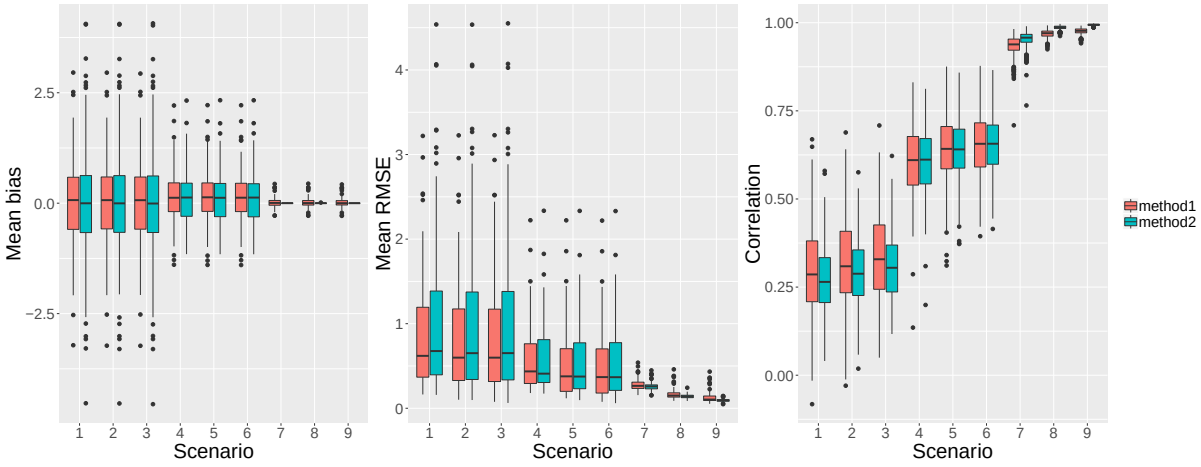

Figure 6: Distribution over areas of the exposure prediction mean bias (left) and mean RMSE (centre), averaged across simulations (see Eq.(10) and (11)), by scenario and upscaling method (Method 1 and Method 2). Right: distribution over simulations of the correlation coefficient between area true exposure values and corresponding predictions, by scenario and upscaling method.

The analysis for the log relative risk parameter γ_1 is based on the results reported in Figure 7. It can be observed that the higher the resolution of the prediction grid (moving from scenarios 1-3 to scenarios 7-9) the lower the parameter bias, independently from the upscaling method. The RMSE index has a slightly different behaviour because the lowest values of the index are observed for scenarios 4-6; in any case all the RMSE medians are small and lower than 0.051. No strong differences are observed across the three Poisson methods (PI, FF and PE) or across upscaling methods (Method 1 and Method 2), with the exception of the bias values for scenarios 1-3 which are always slightly lower for Method 1 even if showing higher variability across simulations. These results are confirmed by the plots of the γ_1 posterior distributions reported in Figure A.1 which show very similar patterns across the 6 cases (3 propagation combined with 2 upscaling methods), in terms of

variability and location, especially for the scenarios from 4 to 9. We expected that the three different methods (PI, FF and PE) employed to acknowledge the exposure uncertainty would have led to differences in the posterior distribution of the Poisson parameter γ_1. This effect is not evident from the plot of Figure A.1 or from the posterior standard deviation (SD) values reported in Table A.2, as the three methods return similar posterior distribution with the same level of precision. Nevertheless, it is worth noting that the upscaling method seems to have an effect on the uncertainty: Method 1 (linear combination with neighbourhood effect) is always associated with a less precise estimation of γ_1 (higher SD), especially for scenarios 1-6.

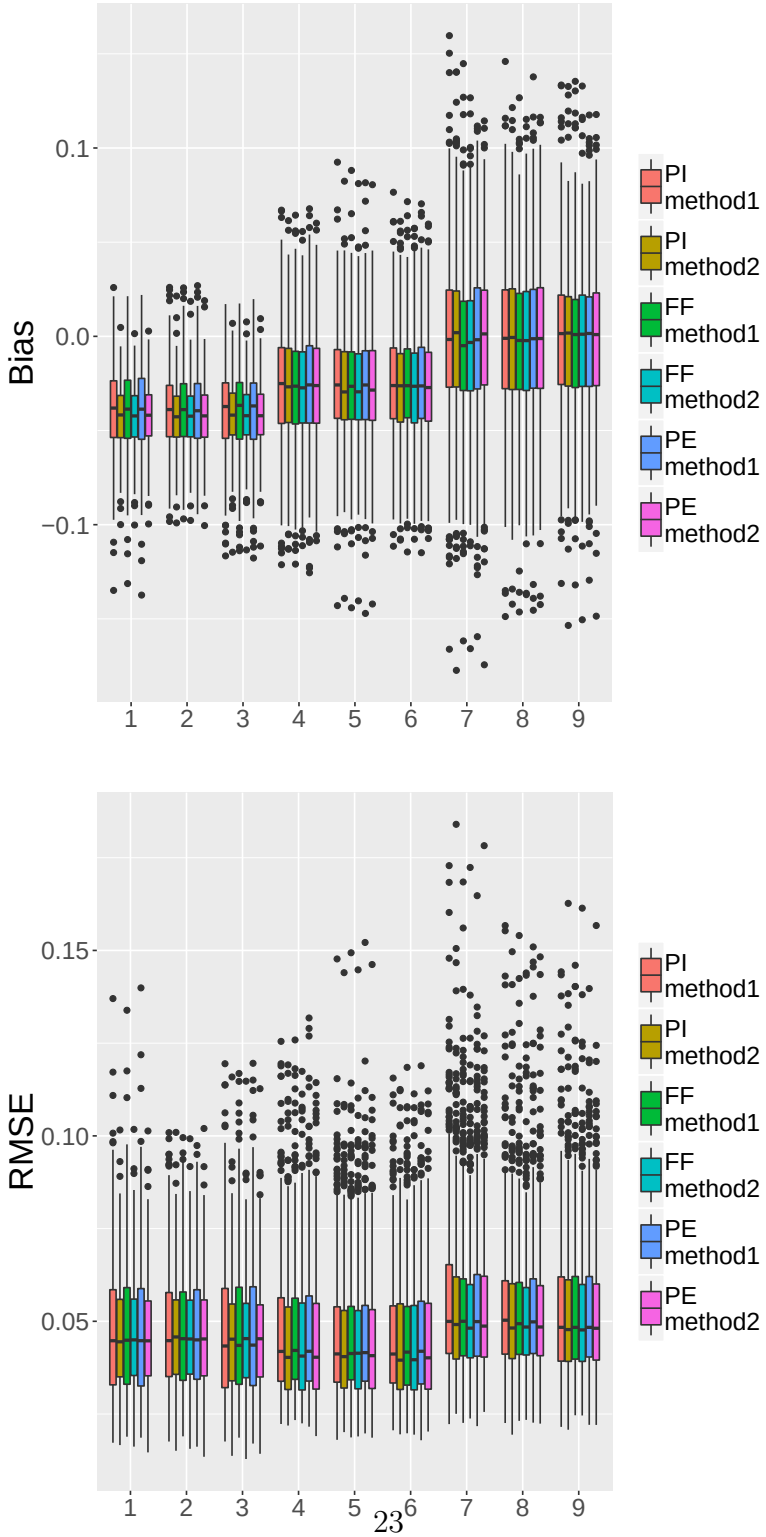

Figure 7: Distribution over simulations of the bias (left) and RMSE (right) for the γ_1 Poisson parameter, by scenario, uncertainty propagation method (PI, FF or PE) and exposure upscaling approach (Method 1, Method 2) for the 9 considered scenarios.

6. Motivating example: results

Going back to the NO_2 and hospitalisations data for the municipalities in the Piemonte region (presented in Section 2), we first estimate the NO_2 exposure model described in Section 3.2 by considering the data from 45 monitoring stations among the 55 available sites (10 sites are randomly chosen and set apart for validation purposes) and a SPDE mesh with 533 vertexes. The vector $\boldsymbol{b}$ of Eq.(5) contains 6 coefficients for the standardized covariates (NO_2 emissions, wind speed, temperature, precipitation, mixing height and altitude). The parameter posterior estimates are reported in Table 2 for Method 1 as upscaling method (the results for Method 2 are not reported because they coincide). It can be observed that NO_2 emissions is the only regressor showing a small but positive posterior mean, with the 95% credible interval completely above zero - while for all the meteorological and geographical variables there is not strong evidence of an effect. The intercept estimate, equal to 3.52 on the log scale, corresponds to a posterior median pollution level of 34.54 $\mu g/m^3$, after adjustment for covariates. The estimate of the measurement error variance σ_e^2 is small (posterior median equal to 0.02) while the variability related to the spatial process is higher with a posterior median for σ_ω^2 equal to 0.10. The posterior mean for σ_ω^2 is 0.12 and this denotes a right skewed posterior distribution; the same happens for the range r which shows a posterior median and mean of about 81 km and 104 km, respectively (consider that the maximum distance in the region is 274 km).

The data of the 10 validation stations are used for computing the prediction performance indexes: the correlation between NO_2 observed and predicted measurements is equal to 0.835, the bias is -0.00003 and the RMSE is equal to 0.231. As a basis for comparison, consider that the universal kriging model with the same covariates and Matèrn variogram performs as follows: correlation 0.732, bias 0.026 and RMSE 0.285. Thus, the exposure model outperforms the universal kriging for NO_2 point prediction.

Figure 8 reports the maps of the posterior medians of NO_2 concentration (transformed back to the original scale in $\mu g/m^3$) for the 56×72 regular grid (left plot) and for the 1206 areas in Piemonte (right plot). As expected, higher concentration are predicted in the areas close to the biggest cities and the main highways that connect Piemonte with Lombardia region on the east. All the municipalities located near the mountain areas surrounding the region on the northern, western and southern side are characterized by

concentration lower than 15 $\mu g/m^3$. These area predictions are obtained using Method 1 for upscaling; Method 2 returns similar patterns, as it can be seen in the difference map reported in Figure A.2. However, we observe that some areas are characterized by different concentration estimates under Method 1 and Method 2 and this may have an effect on the results of the Poisson health model. These differences could be related to the geographical structure of Piemonte region and the number of points used for predicting exposure with Method 1 and Method 2. In this regard, consider that the median (mean) number of prediction points is 4 (5.4) for Method 1 and 1 (1.6) for Method 2 (see also Figure A.3 for the distribution of the number of prediction points across areas).

Parameter	Mean	Sd	0.025quant	Median	0.975quant
Intercept b_0	3.52	0.21	3.06	3.52	3.98
$b_{\text{NO2emissions}}$	0.09	0.04	0.01	0.09	0.17
$b_{\text{windspeed}}$	-0.03	0.05	-0.13	-0.03	0.07
$b_{\text{temperature}}$	-0.06	0.06	-0.18	-0.06	0.06
$b_{\text{precipitation}}$	0.10	0.05	0.00	0.10	0.18
$b_{\text{mixingheight}}$	0.12	0.07	-0.01	0.12	0.25
b_{altitude}	-0.19	0.10	-0.39	-0.19	0.02
σ_e^2	0.02	0.01	0.01	0.02	0.05
σ_ω^2	0.12	0.08	0.04	0.10	0.34
r (in metre)	104249.25	76087.92	31034.37	81308.65	311351.38

Table 2: Posterior summaries (mean, standard deviation (sd), 2.5%, 50% and 97.5% quantiles) of the parameters of the NO_2 exposure model estimated using data from 45 monitoring stations.

The posterior distributions of the Poisson model parameters are reported in Figure 9. No evidence of substantial differences can be seen for the intercept γ_0, the vulnerability parameter γ_2 and the variances σ_v^2 and σ_u^2: all the 6 compared cases (PI, FF and PE with Method 1 and Method 2) provide very similar posterior distributions (see also the posterior summary statistics reported in Table A.3). Some differences can be seen instead for the log-risk parameter γ_1 (for a 10 $\mu g/m^3$ increase in NO_2 concentration), which is the main quantity of interest for assessing the impact of NO_2 exposure on the health outcome. First of all, we can observe that the upscaling method has an effect on the location of the posterior distribution: Method 2 (simple mean of exposure) shifts the distribution to the left (the posterior mean de-

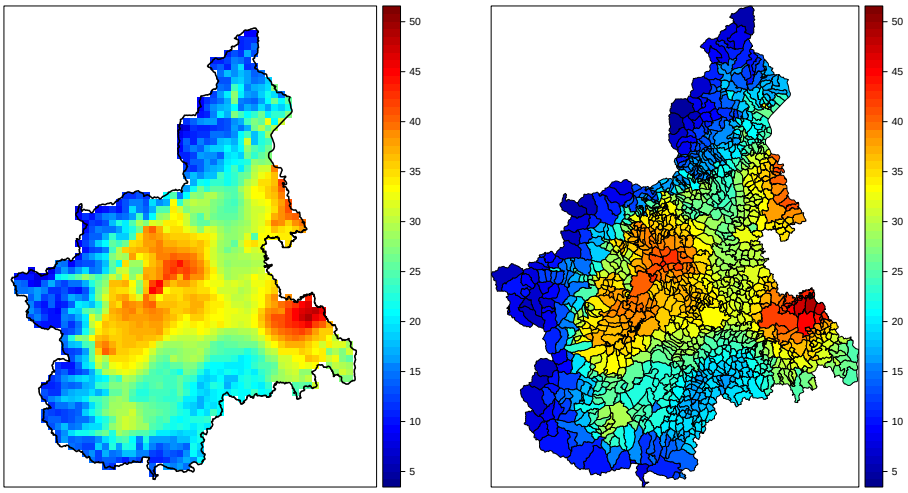

Figure 8: Map of the posterior median of NO_2 concentration (in $\mu g/m^3$) at the grid level (left) and at the area level using Method 1 for upscaling (right).

creases from 0.0016 to 0.0013 for the PI approach, from 0.0109 to 0.0087 for FF and from 0.0177 to 0.0145 for PE) while keeping the variability basically unchanged. The uncertainty propagation method (PI, FF or PE) has an effect both on the location and on the dispersion of the posterior distributions: for the FF approach the γ_1 posterior distribution shows higher variability than the PI and PE strategies, in accordance with the results presented in Blangiardo et al. (2016), and the posterior medians are closer to zero. As a consequence, in the FF case there is not strong evidence that the log-risk parameter γ_1 is different from zero, both for Method 1 and Method 2. The PI approach is the one showing less variability and this is expected since we do not take into account the uncertainty of the exposure estimation. The PE case represents an intermediate situation between PI and FF in terms of variability (for example, considering Method 1, the posterior standard deviation is equal to 0.0072 for PI, 0.0076 for PE and 0.0103 for FF). Finally, note that the 95% credible interval for γ_1 does not include zero only in two cases: PI-Method 1 and PE-Method 1.

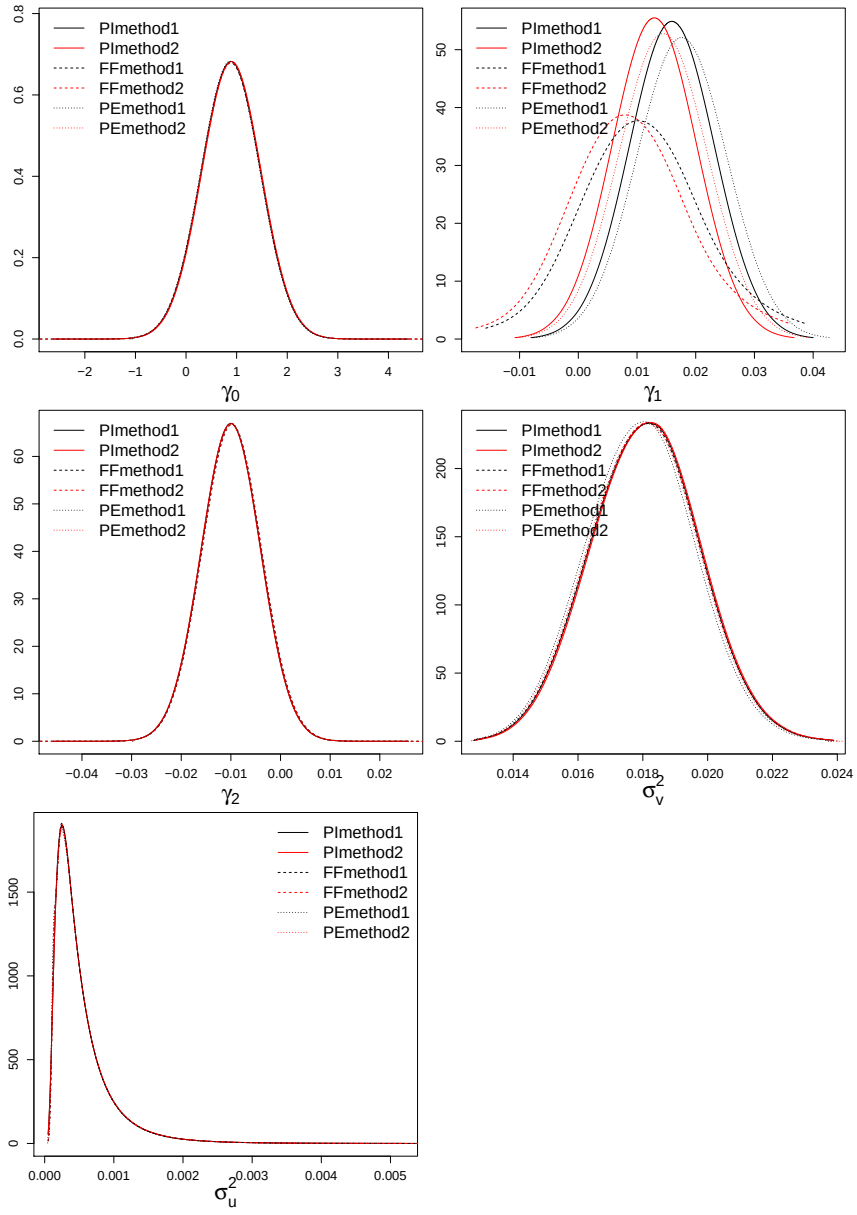

Figure 9: Posterior distributions of the Poisson model parameters: γ_0, γ_1 for a 10 $\mu g/m^3$ increase in NO_2 concentration, the vulnerability index parameter γ_2, the variance σ_v^2 of the iid random effect and the variance σ_u^2 of the spatially structured random effect.

7. Discussion and conclusions

In this paper we presented a two-stage Bayesian model that predict air pollution concentration at the small area level and evaluate its effect on health outcomes. During the first stage the model integrates data from different sources characterized by spatial misalignment: monitoring stations are available at point location, while meteorological variables and NO_2 emissions are available at the monitoring station locations as well as for regular grid points; we have then shown two different methods to aggregate the estimates at the regular grid level to the irregular lattice, which is the spatial resolution available for the health outcome. The second stage links the predicted concentration to the health outcome cases and we have presented different ways of accounting for uncertainty from the first to the second stage. In particular, with the feed-forward approach we draw some samples from the joint NO_2 posterior predictive distributions at the area level and fit for each sample the Poisson health model, thus obtaining a posterior distribution of the risk parameter γ_1. As an alternative, we propose the prior-exposure approach which assumes an informative Normal prior distribution for the area pollutant concentration in the health model (see Eq.(8)). The mean parameters of these priors are spatially correlated as they are taken from the area level exposure predictive distribution estimated by the SPDE approach. Consequently, also the area exposure prior distribution will inherits the spatial structure. Note that instead of using the posterior means and variances from the first stage in the prior, we could sample from the posterior predictive distribution, fit many models with INLA and then obtain a Bayesian model average of the marginals from the fitted models (similarly to Bivand et al., 2014 and Gómez-Rubio and Rue, 2018). As computationally very intensive we have not followed this approach. However, as noted by Gómez-Rubio and Palmí-Perales (2019), fixing some of the parameters to the ML estimates or posterior means may have little impact on the posterior marginals of the remainder of the parameters when fitting (spatial) models with INLA. Another possible extension of the prior-exposure approach of Eq.(8) would consist in using a proper multivariate Normal prior distribution where the mean vector and covariance matrix are determined by posterior predictive distribution samples from the exposure model (see e.g. Warren et al., 2012; Lee et al., 2017).

The main challenge we faced in terms of model evaluation was how to assess how accurate a model predicts area level exposure considering that it is

not possible to perform cross-validation as it is usually done with spatial point prediction since there are no observed values for the pollutant concentration at the area level. In addition, within a simulative approach, it is non trivial to simulate exposure values for administrative regions given that pollutant concentrations is a continuous spatial field which is usually measured in a limited number of monitoring stations. To overcome this issue, in this paper we simulate the *true* exposure for a very big number of spatial points covering completely the considered region (as to recreate a continuous surface) and then we averaged the values across areas.

From the simulation study we found that the exposure performance indexes (bias and RMSE) for the predictions at the area level improves with a higher resolution of the prediction grid. However, it is important to note that commonly in real applications this resolution is fixed and given by the regular grid of the numerical models that provide covariates included in the model. Also the number (and position) of monitoring stations, which has a positive effect on the RMSE and correlation indexes, cannot be chosen by researchers but are set by environmental agencies. To deal with the issue of covariate driven grid size, current work consists in a Bayesian space-time model which integrates several numerical outputs characterised by different spatial resolutions together with ground measurements, and is able to predict pollutant concentration at the desired spatial resolution. With this approach the resolution of the covariate grid does not represent a constraint anymore, as the model reconstructs the spatial fields of the misaligned covariates jointly with the latent field of pollutant concentration from ground measurements in a data assimilation framework.

From the simulation results it can be observed that Method 1 and Method 2 do not seem to have a large effect on the risk parameter γ_1, in terms of bias, RMSE and posterior distribution pattern. However, the results may depend on the spatial variation within and between the areas; as Method 1 averages across more points it should provide more stable estimates of the concentration, particularly in case of local spatial variability. This is more evident on the case study on Piemonte, where Method 2 shifts the posterior distribution of the parameter towards zero. In addition, it is interesting to note that while the FF method of uncertainty propagation results in estimates of the health effect shifted towards 0 and more variable with respect to the other approaches, the PE method leads to higher deviation from 0 with an intermediate level of variability.

The choice of the methods to link stage 1 to stage 2 were motivated by

the fact that we wanted these to be easily adopted by the researchers' wide community; hence we focused on methods which are easily implemented in a readily available R package such as `R-INLA` without the need for algorithm writing. An alternative to the two-stage specification will consist of a joint model of the exposure and of its health effects; this would ensure that uncertainty is directly propagated across it. However the level of computational power required to run the model would increase substantially and at the same time a two-stage approach proves more robust if the exposure model is misspecified, hence we believe that a two-stage approach is to be preferred. Future works will extend the current framework to fit spatio-temporal data and to model multi-pollutant concentrations.

Acknowledgements

MC has been supported by the PRIN EphaStat Project (Project No. 20154X8K23, https://sites.google.com/site/ephastat/) provided by the Italian Ministry for Education, University and Research. MB acknowledges support from the MRC-PHE Centre for Environment and Health, funded by the Medical Research Council (MR/L01341X/1) and Public Health England (PHE). VGR has been supported by grant SBPLY/17/180501/000491, awarded by Consejería de Educación, Cultura y Deportes (JCCM, Spain) and FEDER, and grant MTM2016-77501-P, awarded by Ministerio de Economía y Competitividad (Spain).

Atkinson, R. W., A. Analitis, E. Samoli, G. W. Fuller, D. C. Green, I. S. Mudway, H. R. Anderson, and F. J. Kelly (2016). Short-term exposure to traffic-related air pollution and daily mortality in London, UK. *Journal of Exposure Science and Environmental Epidemiology 26*(2), 125.

Bakar, K. and S. Sahu (2015). sptimer: Spatio-temporal bayesian modeling using r. *Journal of Statistical Software, Articles 63*(15), 1–32.

Bell, M. L. (2006). The use of ambient air quality modeling to estimate individual and population exposure for human health research: A case study of ozone in the Northern Georgia Region of the United States. *Environment International 32*(5), 586–593.

Berrocal, V. J., A. E. Gelfand, and D. M. Holland (2010). A spatio-temporal downscaler for output from numerical models. *Journal of Agricultural, Biological, and Environmental Statistics 15*(2), 176–197.

Besag, J., J. York, and A. Mollié (1991). Bayesian Image Restoration with Two Applications in Spatial Statistics. *The Annals of the Institute of Statistics and Mathematics 43*(1), 1–59.

Bivand, R. and G. Piras (2015). Comparing implementations of estimation methods for spatial econometrics. *Journal of Statistical Software 63*(18), 1–36.

Bivand, R. S., V. Gómez-Rubio, and H. Rue (2014). Approximate Bayesian inference for spatial econometrics models. *Spatial Statistics 9*, 146–165.

Blangiardo, M. and M. Cameletti (2015). *Spatial and Spatio-temporal Bayesian Models with `R-INLA`*. Wiley.

Blangiardo, M., F. Finazzi, and M. Cameletti (2016). Two-stage bayesian model to evaluate the effect of air pollution on chronic respiratory diseases using drug prescriptions. *Spatial and Spatio-temporal Epidemiology 18*, 1–12.

Bruno, F., M. Cameletti, M. Franco-Villoria, F. Greco, R. Ignaccolo, L. Ippoliti, P. Valentini, and M. Ventrucci (2016). A survey on ecological regression for health hazard associated with air pollution. *Spatial Statistics 18*(Part A), 276–299.

Cameletti, M., R. Ignaccolo, and S. Bande (2011). Comparing spatio-temporal models for particulate matter in Piemonte. *Environmetrics 22*(8), 985–996.

Cameletti, M., F. Lindgren, D. Simpson, and H. Rue (2013). Spatio-temporal modeling of particulate matter concentration through the SPDE approach. *AStA Advances in Statistical Analysis 97*(2), 109–131.

Carey, I. M., H. R. Anderson, R. W. Atkinson, S. Beevers, D. G. Cook, D. Dajnak, J. Gulliver, and F. J. Kelly (2016). Traffic pollution and the incidence of cardiorespiratory outcomes in an adult cohort in London. *Occupational and Environmental Medicine 73*(12), 849–856.

Carugno, M., D. Consonni, G. Randi, D. Catelan, L. Grisotto, P. A. Bertazzi, A. Biggeri, and M. Baccini (2016). Air pollution exposure, cause-specific deaths and hospitalizations in a highly polluted Italian region. *Environmental research 147*, 415–424.

EEA (2017a). Air quality in Europe. Technical report, European Environmental Agency (EEA). https://www.eea.europa.eu/publications/air-quality-in-europe-2017/at_download/file.

EEA (2017b). Urban PM2.5 Atlas: Air Quality in European cities. Technical report, European Environmental Agency (EEA). https://ec.europa.eu/jrc/en/publication/eur-scientific-and-technical-research-reports/urban-pm25-atlas-air-quality-european-cities.

Elliott, P., G. Shaddick, J. C. Wakefield, C. d. Hoogh, and D. J. Briggs (2007). Long-term associations of outdoor air pollution with mortality in Great Britain. *Thorax 62*(12), 1088–1094.

Fassò, A. and F. Finazzi (2011). Maximum likelihood estimation of the dynamic coregionalization model with heterotopic data. *Environmetrics 22*(6), 735–748.

Faustini, A., R. Rapp, and F. Forastiere (2014). Nitrogen dioxide and mortality: review and meta-analysis of long-term studies. *European Respiratory Journal 44*(3), 744–753.

Fuentes, M., H.-R. Song, S. K. Ghosh, D. M. Holland, and J. M. Davis (2006). Spatial association between speciated fine particles and mortality. *Biometrics 62*(3), 855–863.

Gelfand, A. (2010). *Handbook of Spatial Statistics*, Chapter Misaligned Spatial Data: The Change of Support Problem. Chapman & Hall.

Gómez-Rubio, V. and F. Palmí-Perales (2019). Multivariate posterior inference for spatial models with the integrated nested laplace approximation. *Journal of the Royal Statistical Society: Series C (Applied Statistics) 68*(1), 199–215.

Gómez-Rubio, V. and H. Rue (2018). Markov chain Monte Carlo with the Integrated Nested Laplace Approximation. *Statistics and Computing 28*(5), 1033–1051.

Halonen, J. I., M. Blangiardo, M. B. Toledano, D. Fecht, J. Gulliver, H. R. Anderson, S. D. Beevers, D. Dajnak, F. J. Kelly, and C. Tonne (2016).

Long-term exposure to traffic pollution and hospital admissions in London. *Environmental Pollution 208*, 48–57.

Halonen, J. I., M. Blangiardo, M. B. Toledano, D. Fecht, J. Gulliver, R. Ghosh, H. R. Anderson, S. D. Beevers, D. Dajnak, F. J. Kelly, et al. (2016). Is long-term exposure to traffic pollution associated with mortality? A small-area study in London. *Environmental Pollution 208*, 25–32.

Huang, G., D. Lee, and M. Scott (2015). An integrated Bayesian model for estimating the long-term health effects of air pollution by fusing modelled and measured pollution data: a case study of nitrogen dioxide concentrations in Scotland. *Spatial and Spatio-temporal Epidemiology 14-15*, 63–74.

Huang, G., M. Scott, and D. Lee (2017). Multivariate space-time modelling of multiple air pollutants and their health effects accounting for exposure uncertainty. *Statistics in Medicine*, 1–15.

Lee, A., A. Szpiro, S. Kim, and L. Sheppard (2015). Impact of preferential sampling on exposure prediction and health effect inference in the context of air pollution epidemiology. *Environmetrics 26*(4), 255–267.

Lee, D. (2011). A comparison of conditional autoregressive models used in Bayesian disease mapping. *Spatial and Spatio-Temporal Epidemiology 2*(2), 79–89.

Lee, D. (2018). A locally adaptive process-convolution model for estimating the health impact of air pollution. *Annals of Applied Statistics*.

Lee, D., S. Mukhopadhyay, A. Rushworth, and S. K. Sahu (2017). A rigorous statistical framework for spatio-temporal pollution prediction and estimation of its long-term impact on health. *Biostatistics 18*(2), 370–385.

Lee, D. and S. K. Sahu (2016). *Handbook of Spatial Epidemiology*, Chapter Estimating the Health Impact of Air Pollution Fields. Chapman & Hall.

Lee, D. and C. Sarran (2015). Controlling for unmeasured confounding and spatial misalignment in long-term air pollution and health studies. *Environmetrics 26*, 477–487.

Lee, D. and G. Shaddick (2010). Spatial modeling of air pollution in studies of its short-term health effects. *Biometrics 66*(4), 1238–1246.

Lim, S. S., T. Vos, et al. (2012). A comparative risk assessment of burden of disease and injury attributable to 67 risk factors and risk factor clusters in 21 regions, 1990-2010: a systematic analysis for the global burden of disease study 2010. *The Lancet 380*(9859), 2224–2260.

Lindgren, F. and H. Rue (2015). Bayesian Spatial Modelling with `R-INLA`. *Journal of Statistical Software 63*(19), 1–25.

Lindgren, F., H. Rue, and J. Lindström (2011). An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach. *Journal of the Royal Statistical Society: Series B 73*(4), 423–498.

Liu, Y., G. Shaddick, and J. V. Zidek (2016). Incorporating high-dimensional exposure modelling into studies of air pollution and health. *Statistics in Biosciences*, 1–23.

Madsen, L., D. Ruppert, and N. S. Altman (2008). Regression with spatially misaligned data. *Environmetrics 19*(5), 453–467.

Moore, E., L. Chatzidiakou, M.-O. Kuku, R. L. Jones, L. Smeeth, S. Beevers, F. J. Kelly, B. Barratt, and J. K. Quint (2016). Global associations between air pollutants and chronic obstructive pulmonary disease hospitalizations. a systematic review. *Annals of the American Thoracic Society 13*(10), 1814–1827.

Moraga, P., S. M. Cramb, K. L. Mengersen, and M. Pagano (2017). A geostatistical model for combined analysis of point-level and area-level data using inla and spde. *Spatial Statistics 21*, 27 – 41.

OECD (2016). The economic consequences of outdoor air pollution. Technical report, Organisation for Economic Co-operation and Development (OECD). http://www.oecd.org/environment/indicators-modelling-outlooks/Policy-Highlights-Economic-consequences-of-outdoor-air-pollution-web.pdf.

Pannullo, F., D. Lee, E. Waclawski, and A. H. Leyland (2015). Improving spatial nitrogen dioxide prediction using diffusion tubes: A case study in West Central Scotland. *Atmospheric Environment 118*, 227–235.

Pannullo, F., D. Lee, E. Waclawski, and A. H. Leyland (2016). How robust are the estimated effects of air pollution on health? accounting for model uncertainty using bayesian model averaging. *Spatial and Spatio-temporal Epidemiology 18*, 53–62.

Peng, R. D. and M. L. Bell (2010). Spatial misalignment in time series studies of air pollution and health data. *Biostatistics (Oxford, England) 11*(4), 720–740.

Pirani, M., J. Gulliver, G. W. Fuller, and M. Blangiardo (2014). Bayesian spatiotemporal modelling for the assessment of short-term exposure to particle pollution in urban areas. *Journal of Exposure Science and Environmental Epidemiology 24*(3), 319–327.

Powell, H. and D. Lee (2014). Modelling spatial variability in concentrations of single pollutants and composite air quality indicators in health effects studies. *Journal of the Royal Statistical Society: Series A (Statistics in Society) 177*(3), 607–623.

Raaschou-Nielsen, O., Z. J. Andersen, S. S. Jensen, M. Ketzel, M. Sørensen, J. Hansen, S. Loft, A. Tjønneland, and K. Overvad (2012). Traffic air pollution and mortality from cardiovascular disease and all causes: a danish cohort study. *Environmental Health 11*(1), 60.

Rue, H., S. Martino, and N. Chopin (2009). Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations. *Journal of the Royal Statistical Society: Series B 2*(71), 1–35.

Rue, H., A. Riebler, S. H. S., J. B. Illian, D. P. Simpson, and F. K. Lindgren (2017). Bayesian computing with inla: A review. *Annual Review of Statistics and Its Application 4*(1), 395–421.

Rushworth, A., D. Lee, and R. Mitchell (2014). A spatio-temporal model for estimating the long-term effects of air pollution on respiratory hospital admissions in Greater London. *Spatial and spatio-temporal epidemiology 10*, 29–38.

Sahu, S. (2011). Hierarchical Bayesian models for space-time air pollution data. In C. Rao (Ed.), *Handbook of Statistics-Time Series Analysis, Methods and Applications*, Volume 30 of *Handbook of Statistics*. Elsevier Publishers, Holland.

Sahu, S., A. Gelfand, and D. Holland (2010). Fusing point and areal level spacetime data with application to wet deposition. *Journal of the Royal Statistical Society: Series C 59*(1), 77–103.

Sanyal, S., T. Rochereau, C. N. Maesano, L. Com-Ruelle, and I. Annesi-Maesano (2018, 11). Long-term effect of outdoor air pollution on mortality and morbidity: A 12-year follow-up study for metropolitan france. *International journal of environmental research and public health 15*(11), 2487.

Shaddick, G., M. L. Thomas, A. Green, M. Brauer, A. van Donkelaar, R. Burnett, H. H. Chang, A. Cohen, R. V. Dingenen, C. Dora, S. Gumy, Y. Liu, R. Martin, L. A. Waller, J. West, J. V. Zidek, and A. Prüss-Ustün (2018). Data integration model for air quality: a hierarchical approach to the global estimation of exposures to ambient air pollution. *Journal of the Royal Statistical Society: Series C (Applied Statistics) 67*(1), 231–253.

Shaddick, G. and J. V. Zidek (2016). *Spatio-Temporal Methods in Environmental Epidemiology*. CRC Press.

Wakefield, J. and G. Shaddick (2006). Health-exposure modeling and the ecological fallacy. *Biostatistics 7*(3), 438–455.

Wang, Y., M. Pirani, A. L. Hansell, S. Richardson, and M. Blangiardo (2019). Using ecological propensity score to adjust for missing confounders in small area studies. *Biostatistics (Oxford, England) 20*(1), 1–16.

Warren, J., M. Fuentes, A. Herring, and P. Langlois (2012, 12). Spatial-temporal modeling of the association between air pollution exposure and preterm birth: identifying critical windows of exposure. *Biometrics 68*(4), 1157–1167.

Young, L. J., C. A. Gotway, J. Yang, G. Kearney, and C. DuClos (2009). Linking health and environmental data in geographical analysis: It's so much more than centroids. *Spatial and Spatio-temporal Epidemiology 1*(1), 73–84.

Zhu, L., B. P. Carlin, and A. Gelfand (2003). Hierarchical regression with misaligned spatial data: relating ambient ozone and pediatric asthma ER visits in Atlanta. *Environmetrics 14*(5), 537–557.

Appendix A. Extra tables and figures

	b_0		b_1		σ_e^2		σ_ξ^2		r	
Scenario	Bias	RMSE	Bias	RMSE	Bias	RMSE	Bias	RMSE	Bias	RMSE
1	-0.044	1.033	-0.001	0.064	-0.012	0.039	0.181	0.708	0.188	1.450
2	0.001	0.845	-0.001	0.023	0.001	0.012	0.151	0.624	0.174	1.007
3	-0.018	0.830	-0.000	0.013	0.001	0.006	0.195	0.683	0.195	0.869
4	-0.044	1.032	-0.001	0.064	-0.012	0.039	0.180	0.707	0.188	1.450
5	-0.014	0.838	-0.001	0.023	0.001	0.011	0.155	0.638	0.186	0.997
6	-0.012	0.821	0.000	0.013	0.001	0.006	0.204	0.688	0.200	0.903
7	-0.044	1.039	-0.001	0.064	-0.012	0.040	0.181	0.708	0.196	1.461
8	-0.014	0.838	-0.001	0.023	0.001	0.011	0.154	0.636	0.185	0.996
9	-0.012	0.822	0.000	0.013	0.001	0.006	0.206	0.692	0.200	0.903

Table A.1: Bias and RMSE (averaged across simulations) for the spatial parameters.

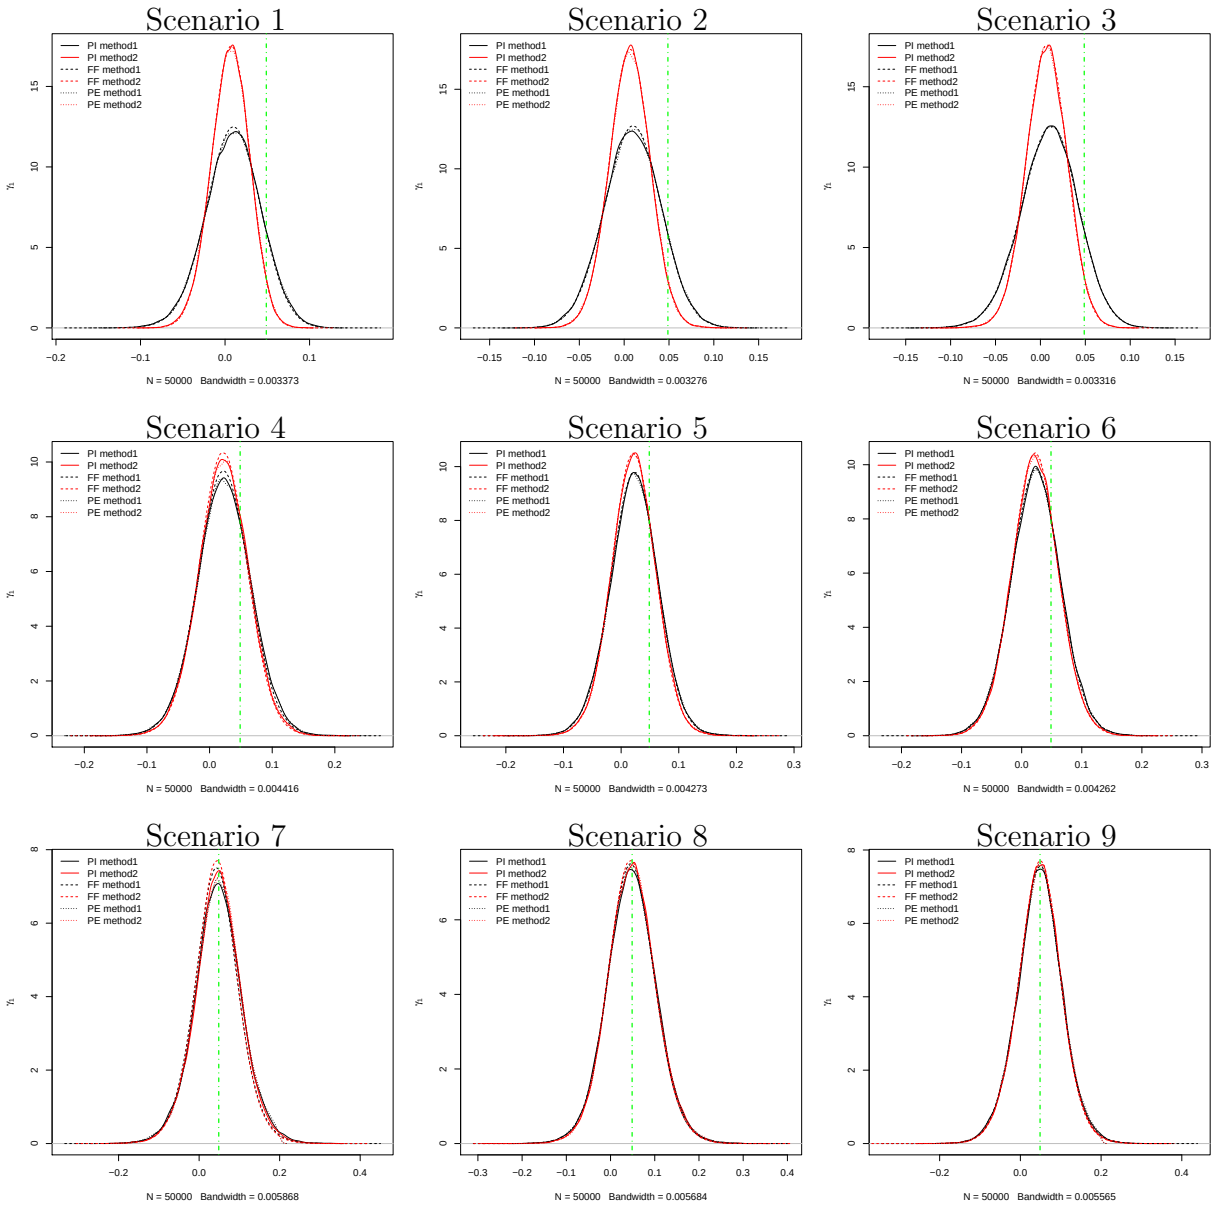

Figure A.1: Posterior distributions of the Poisson model parameter γ_1 for the 9 scenarios considered in the simulation study according to the upscaling method (Method 1, Method 2) and the uncertainty propagation approach (PI, FF, PE). The vertical green line represents the true parameter value.

Scenario	Propagation method	Upscaling method	
		Method 1	Method 2
1	PI	0.033	0.023
1	FF	0.032	0.023
1	PE	0.033	0.023
2	PI	0.032	0.023
2	FF	0.032	0.023
2	PE	0.032	0.023
3	PI	0.032	0.023
3	FF	0.032	0.023
3	PE	0.032	0.023
4	PI	0.043	0.040
4	FF	0.042	0.039
4	PE	0.043	0.040
5	PI	0.042	0.039
5	FF	0.042	0.039
5	PE	0.042	0.039
6	PI	0.042	0.039
6	FF	0.042	0.039
6	PE	0.042	0.040
7	PI	0.060	0.058
7	FF	0.058	0.056
7	PE	0.058	0.057
8	PI	0.057	0.055
8	FF	0.056	0.055
8	PE	0.057	0.055
9	PI	0.057	0.055
9	FF	0.056	0.055
9	PE	0.056	0.054

Table A.2: Posterior standard deviation for the Poisson model parameter γ_1 for the 9 scenarios considered in the simulation study according to the upscaling method (Method 1, Method 2) and the uncertainty propagation approach (PI, FF, PE).

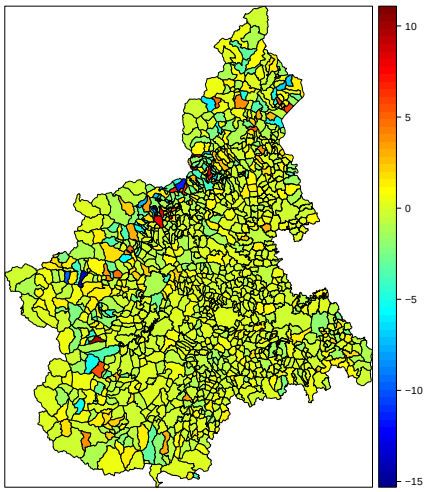

Figure A.2: Map of the differences between the posterior medians of NO_2 concentrations (in $\mu g/m^3$) obtained using Method 1 and Method 2.

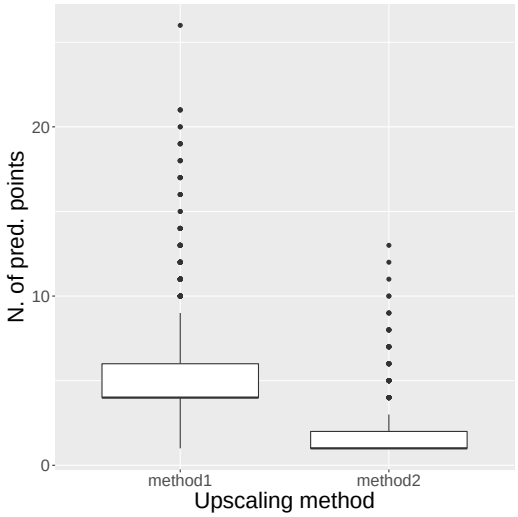

Figure A.3: Distribution of the number of prediction points across areas of the Piemonte region according to the upscaling method (Method 1 and Method 2).

Parameter	Method	Mean	Sd	Quant0.025	Median	Quant0.975
	PI method1	0.8866	0.5850	-0.2622	0.8851	2.0308
	PI method2	0.8989	0.5851	-0.2500	0.8974	2.0433
γ_0	FF method1	0.8824	0.5879	-0.2738	0.8798	2.0309
	FF method2	0.8916	0.5873	-0.2633	0.8890	2.0391
	PE method1	0.8741	0.5817	-0.2693	0.8733	2.0147
	PE method1	0.8887	0.5817	-0.2551	0.8880	2.0290
	PI method1	0.0159	0.0072	0.0018	0.0159	0.0301
	PI method2	0.0130	0.0072	-0.0011	0.0129	0.0270
γ_1	FF method1	0.0109	0.0103	-0.0088	0.0106	0.0326
	FF method2	0.0087	0.0100	-0.0105	0.0084	0.0298
	PE method1	0.0177	0.0076	0.0027	0.0176	0.0326
	PE method2	0.0145	0.0075	-0.0003	0.0145	0.0293
	PI method1	-0.0100	0.0060	-0.0217	-0.0100	0.0017
	PI method2	-0.0100	0.0060	-0.0217	-0.0101	0.0016
γ_2	FF method1	-0.0098	0.0060	-0.0216	-0.0098	0.0019
	FF method2	-0.0098	0.0060	-0.0216	-0.0099	0.0019
	PE method1	-0.0099	0.0059	-0.0216	-0.0099	0.0017
	PE method2	-0.0100	0.0059	-0.0216	-0.0100	0.0016
	PI method1	0.0181	0.0017	0.0148	0.0181	0.0214
	PI method2	0.0181	0.0017	0.0149	0.0181	0.0214
σ_v^2	FF method1	0.0181	0.0017	0.0148	0.0181	0.0214
	FF method2	0.0181	0.0017	0.0148	0.0181	0.0214
	PE method1	0.0179	0.0017	0.0147	0.0179	0.0212
	PE method2	0.0180	0.0017	0.0148	0.0180	0.0213
	PI method1	0.0005	0.0004	0.0001	0.0004	0.0017
	PI method2	0.0005	0.0004	0.0001	0.0004	0.0017
σ_u^2	FF method1	0.0005	0.0004	0.0001	0.0004	0.0017
	FF method2	0.0005	0.0004	0.0001	0.0004	0.0017
	PE method1	0.0005	0.0004	0.0001	0.0004	0.0017
	PE method2	0.0005	0.0004	0.0001	0.0004	0.0017

Table A.3: Posterior summary statistics (mean, standard deviation (sd), 2.5%, 50% and 97.5% quantiles) for the Poisson model parameters according to the upscaling method (Method 1, Method2) and the propagation approach (PI, FF, PE). Recall that γ_1 refers to a 10 $\mu g/m^3$ increase in NO_2 concentration.