



Published in final edited form as:

Ann Appl Stat. 2010 December 1; 4(4): 1942–1975.

A bivariate space-time downscaler under space and time misalignment

Veronica J. Berrocal, Alan E. Gelfand, and David M. Holland¹

Abstract

Ozone and particulate matter $PM_{2.5}$ are co-pollutants that have long been associated with increased public health risks. Information on concentration levels for both pollutants come from two sources: monitoring sites and output from complex numerical models that produce concentration surfaces over large spatial regions. In this paper, we offer a fully-model based approach for fusing these two sources of information for the pair of co-pollutants which is computationally feasible over large spatial regions and long periods of time. Due to the association between concentration levels of the two environmental contaminants, it is expected that information regarding one will help to improve prediction of the other. Misalignment is an obvious issue since the monitoring networks for the two contaminants only partly intersect and because the collection rate for $PM_{2.5}$ is typically less frequent than that for ozone.

Extending previous work in Berrocal et al. (2009), we introduce a bivariate downscaler that provides a flexible class of bivariate space-time assimilation models. We discuss computational issues for model fitting and analyze a dataset for ozone and $PM_{2.5}$ for the ozone season during year 2002. We show a modest improvement in predictive performance, not surprising in a setting where we can anticipate only a small gain.

Keywords

Co-kriging; coregionalization; dynamic model; kriging; multivariate spatial process; spatially varying coefficients

1 Introduction

Ozone and particulate matter, $PM_{2.5}$, have long been associated with increased public health risks, e.g., of respiratory diseases (Schwartz 1996; Dominici et al. 2006; Braga et al. 2001), cardiovascular diseases (Dominici et al. 2006; Braga et al. 2001), and mortality and morbidity in general (Dominici et al. 2000; Smith et al. 2000). To set air quality standards the EPA utilizes information from monitoring networks and from air quality computer models.

The two sources of information are valuable in different ways. The observations reported by monitoring networks, though sparsely collected and, sometimes, affected by missingness, provide direct measurements of the true value up to measurement error. The output from air

¹Veronica J. Berrocal is Postdoctoral Fellow, Statistical and Applied Mathematical Science Institute, Research Triangle Park, NC 27709, USA (vjb2@stat.duke.edu). Alan E. Gelfand is Professor, Department of Statistical Science, Duke University, Durham, NC 27708, USA (alan@stat.duke.edu). David M. Holland is senior statistician, U.S. Environmental Protection Agency, National Exposure Research Laboratory, Research Triangle Park, NC 27711, USA (Holland.David@epa.gov).

Disclaimer: The U.S. Environmental Protection Agency through its Office of Research and Development partially collaborated in the research described here. Although it has been reviewed by the Agency and approved for publication, it does not necessarily reflect the Agency's policies or views.

quality models estimates pollutant concentrations over large spatial domains at the grid cell level. The estimates are viewed as averages over these cells and do not contain any missing data, but are uncalibrated. It is desirable to combine both sources of information but to do so, it is necessary to solve the inherent “change of support” problem of the data (Cressie 1993; Gotway and Young 2002; Banerjee et al. 2004), that is, the spatial misalignment between the observational data and the numerical model output. Addressing the difference in spatial resolution between the two sources of data will also allow to evaluate and calibrate the numerical model output.

We are not interested in calibration of the numerical model in the sense of Kennedy and O’Hagan (2001). Rather our focus is on resolving the difference in spatial scale between the numerical model output and the observational data, while bias-correcting the predictions generated by a numerical model. In this regard, several methods have been proposed to assess numerical models and address the issue of “incommensurability” between grid model averages and point measurements (Swall and Foley 2009). Such effort is also common in the context of data assimilation within the atmospheric science literature, where the goal is to combine observational data on the current state of the atmosphere with a short-range forecast produced by a numerical weather prediction model to yield initial conditions for a numerical atmospheric model. The approaches tend to be algorithmic and ad hoc, only occasionally model-based, and do not necessarily address downscaling. Fuller discussion can be found in Daley (1991) and Kalnay (2003). In the context of air quality, Meiring et al. (1998) propose to evaluate the hourly ozone predictions generated by a geophysical model on the grid-cell scale by using the observations taken at monitoring sites, and predicting hourly level of ozone at the model grid cells by averaging the predictions at M regularly-spaced sites taken within each grid cell (see, as well, Fuentes et al. (2003)). A different strategy has been proposed by Jun and Stein (2004) who ignore the difference in spatial resolution between model output and observations, rather suggesting to evaluate a numerical model by looking for differences between the model output and the observations in terms of variograms and correlograms.

Fuentes and Raftery (2005) develop a “fusion” model, expressing the numerical model output as the integral over a grid cell (scaled by the area of the cell) of a latent point-level process. Their main goal is to recover the true unobserved process, but, as a by-product, they obtain estimates of the bias of the numerical model output, thus allowing evaluation and calibration of the numerical model. The work is an instance of Bayesian melding (Poole and Raftery 2000), and has gained considerable popularity (Smith and Cowles 2007; Foley and Fuentes 2008). See also recent work in this spirit applied to sulfate aerosol (Swall and Davis 2006) and ammonium (Davis and Swall 2006). The Bayesian melding approach of Fuentes and Raftery (2005), though popular, suffers from several limitations, as pointed out by Liu et al. (2007) and Berrocal et al. (2009). Additionally, it is only computationally feasible as a spatial model. A spatio-temporal extension, proposed by McMillan et al. (2009), specifies the latent process at the grid cell level, leading to comparison between the numerical model output and the observational data at the grid cell scale.

Operating at the point level, Guillas et al. (2008) propose to correct the model output by downscaling the predictions from grid cells to points and comparing them with the observations. For each site, a two-step linear regression model that relates the observed hourly ozone level at a site with the numerical model output at the grid cell in which the site lies is proposed. Building upon the work of Guillas et al. (2008), Liu et al. (2008) suggest interpolating the intercept and the coefficient of the numerical model output, estimated at each site, via kriging, thus developing a “spatio-temporal” model that will allow *corrected* prediction of ozone level also at unmonitored sites. However, this ad-hoc approach does not provide estimates of the uncertainties associated with the predictions. A formal spatio-

temporal model, that builds upon the downscaling idea of Guillas et al. (2008), has been proposed as a univariate downscaler by Berrocal et al. (2009) with a version implemented and described in Sahu et al. (2009) and in Gelfand and Sahu (2009).

The contribution of this paper is to extend the spatio-temporal downscaler from the univariate setting to a bivariate setting in order to exploit the correlation both in the observed concentration levels of ozone and $PM_{2.5}$ and in the concentration levels of ozone and $PM_{2.5}$, provided by the numerical model. Working in a bivariate setting will prove particularly advantageous for predicting particulate matter. The sampling frequency of the $PM_{2.5}$ monitors - most of the monitors measure $PM_{2.5}$ concentration every 3 days - yields challenging interpolation of the monitoring data to the entire spatial domain as well as difficult evaluation of the predictions generated by the numerical model. Rather, in environmental health studies, researchers always use monitoring data to characterize exposure to fine particulate matter, dealing with the $PM_{2.5}$ sampling frequency issue by aggregating the monitoring data over time and/or over space. Similarly, most of the research effort on $PM_{2.5}$ has been devoted to developing spatial and spatio-temporal models to predict $PM_{2.5}$ based on monitoring data, meteorological covariates or observed concentrations for another pollutant (Brown et al. 1994; Le et al. 1997; Kibria et al. 2002; Smith et al. 2003; Sahu and Mardia 2005; Sahu et al. 2006). A fusion model to combine monitoring data for $PM_{2.5}$ with satellite aerosol optical depth (AOD) data has been proposed by Paciorek and Liu (2008): however, the analysis was conducted on monthly average concentration rather than daily and it revealed that AOD provides no improvement in predicting fine particulate matter.

In this paper, we present a general bivariate spatio-temporal downscaler developed for the numerical model outputs of an air quality model; we illustrate the methodology with regard to prediction of concentration of ozone and fine particulate matter. The model, which to our knowledge, is the first bivariate fusion/downscaler model, extends the univariate downscaler of Berrocal et al. (2009), and, following Berrocal et al. (2009), regresses the bivariate vector of observed ozone and $PM_{2.5}$ concentration on the numerical model outputs for both pollutants using spatially-varying coefficients (Schmidt and Gelfand 2003; Gelfand et al. 2004), modeled as a correlated six-dimensional Gaussian process. The model specification is general and flexible, while offering feasible computation for a fully model-based fusion across space and time. We explore several different special cases, corresponding to different assumptions on the correlation structure of the two pollutants. An important feature of our model is that it allows us to handle not only the spatial misalignment between monitoring data and model output, but also it allows us to accommodate the spatial and temporal misalignment between the ozone and $PM_{2.5}$ monitoring data. Finally, exploiting the correlation, not only between the model output and the corresponding monitoring data for a pollutant, but also between pollutants, we can jointly predict ozone and $PM_{2.5}$ at any site in the spatial domain and provide a measure of the uncertainties associated with such predictions. Additionally, at sites where one pollutant is measured but the other is not, we can predict the level of the latter.

The paper is organized as follows. In Section 2, we present the monitoring data and the numerical model output, in Section 3 we describe the bivariate downscaler model, by first introducing the downscaler in the univariate setting, and then extending the model to the bivariate case, first in a purely spatial setting and then in a spatio-temporal setting. Details on how to handle the spatio-temporal misalignment in the monitoring data are also discussed in Section 3. Section 4 presents details on the model fitting, while Section 5 displays results of our analysis. Finally, we conclude in Section 6 with a discussion and evaluation of our method and with indications for future work.

2 Data

Particulate matter (PM_{2.5}) and ozone are two of the “criteria pollutants” that the Environmental Protection Agency (EPA) is required to monitor by the Clean Air Act. The first is a mixture of solid and liquid particles, emitted in the atmosphere either directly from a source or as result of complicated reactions of chemicals, while the second is a gas made of three atoms of oxygen that is created by a chemical reaction between oxides of nitrogen and volatile organic compounds in the presence of sunlight.

The EPA tracks both criteria pollutants using both measurements taken at monitoring sites and estimates of air pollutants concentration produced by the Models-3/Community Mesoscale Air Quality (CMAQ) model (Byun and Schere 2006; epa.gov/asmdnerl/CMAQ). The latter is a deterministic numerical model that predicts concentrations for various pollutants by integrating three major components: an atmospheric component accounting for the atmosphere and its states and motions, an emission component accounting for the emissions injected in the atmosphere, and a chemical component accounting for the reactions between the different gas present in the atmosphere. By simulating various chemical and physical processes, such as horizontal and vertical advection, emission injection, deposition, plume chemistry effects, etc., CMAQ produces estimates of pollutants concentration at pre-determined spatial scales, e.g., 36 km, 12 km, and, most recently, at 4 km grid cells.

We consider daily estimates and measurements of ozone and PM_{2.5} concentration for the region in the South Eastern part of the United States shown in Figure 1 during the period June 1 - September 30, 2002, the summer season when solar radiation and temperature are high. In turn this encourages not only high concentrations of ozone but also of particulate SO₄, the dominant component of PM_{2.5} in the eastern U.S as well as particulate ammonium nitrate. Relatively strong association between ozone and PM_{2.5} concentrations is expected and, in fact, observed. Following the air quality standards (NAAQS; <http://www.epa.gov/air/criteria.html>) set by EPA for ozone and PM_{2.5}, daily ozone concentration is measured as the daily maximum 8-hour average concentration, while daily PM_{2.5} concentration is given by the 24-hour average concentration of particulate matter.

The monitoring data used in our case study comes from 226 sites sparsely located in the region, and it has been obtained from the EPA Air Quality System (AQS) repository database. Of these 226 sites, not all measure both pollutants: 71 report only ozone measurements, 50 report only PM_{2.5}, and 105 measure both, pinpointing the spatial misalignment in the monitoring data. In order to validate the model with out-of-sample predictions, we randomly split the sites into two sets: a training set used to fit the model, and a validation set used to assess the performance of the model. In the training set, 52 sites measured only ozone concentration, 39 reported only PM_{2.5} concentration, and 70 reported both. In total, the dataset used to fit the model comprised 14,630 *daily* measurements of ozone concentration and 4,790 measurements of PM_{2.5} concentration. The smaller number of observations for particulate matter is due to the sampling scheme for PM_{2.5}: of the 109 PM_{2.5} monitoring sites, only 11 (10.1%) measure concentration of particulate matter every day, while 90 (82.5%) report measurements every three days, and 8 (7.3%) measure PM_{2.5} every six days.

The difference in sampling frequency between ozone and PM_{2.5} suggests that, in addition to spatial misalignment, there is also temporal misalignment in the monitoring data: even sites that measure both pollutants do not sample the two pollutants with the same temporal frequency. However, when concentration measurements for both pollutant are available, they display considerable correlation, equal to 0.69. Spatial and temporal misalignment is

also present among the validation sites: 54 measured only ozone, 46 only $PM_{2.5}$ sites, while 35 reported concentration for both pollutants. In total, the validation dataset contained 6,530 daily observations of ozone concentration and 2,559 daily observations of concentration of particulate matter. The location of the sites used to fit the and validate the model is shown in Figure 1.

An exploratory analysis of the observed daily concentration data for ozone and $PM_{2.5}$ revealed the need for a transformation. We modeled the daily ozone concentration data on the square root scale to stabilize the variance, while we log-transformed the $PM_{2.5}$ concentration data to achieve normality. Both transformations have been previously used; see, e.g., Sahu et al. 2007, Sahu et al. 2006, Carroll et al. 1997, Haslett and Raftery 1989. A normal Q-Q plot of the square root of the observed daily concentration of ozone at sites used to fit the model is shown in Figure 2 (a), while Figure 2 (b) presents a normal Q-Q plot of the logarithm of daily average concentration of $PM_{2.5}$ at the 109 monitoring sites in the test set measuring particulate matter. Both plots seems to suggest a slight deviation from normality, however the deviation is minimal and we are comfortable using a normal model, in agreement with the literature. Nevertheless, in Section 6 we address distributional issues more at lengths and suggest an extension of our downscaling approach for non-Gaussian variables. The overall mean and standard deviation for the square root of ozone concentration and for the log of $PM_{2.5}$ concentration at the test sites were, respectively, 7.41 and $1.31 \sqrt{\text{ppb}}$ (ppb:parts per billions), and 2.72 and $0.56 \log(\mu\text{g}/\text{m}^3)$.

Estimates of the daily average and standard deviation for ozone and $PM_{2.5}$ are presented in Figure 3. Both panels in Figure 3 indicate daily variability in the concentration level for the pollutants, with Figure 3(a) revealing the seasonality of ozone during the study period (June 1 - September 20, 2002). In terms of variability, both ozone and $PM_{2.5}$ present the largest standard deviation on June 25, 2002. For this day, plots of the observed ozone and $PM_{2.5}$ concentration are shown, on the original scales, respectively, in Figure 5(a) and Figure 6(a). (These plots are developed in conjunction with the data analysis of Section 5.)

The numerical model output data consists of estimates of ozone and $PM_{2.5}$ concentration level generated by the CMAQ numerical model run at 12 km grid cell resolution. An example of the type of concentration surfaces yielded by CMAQ can be observed, respectively, in Figure 5(b) and Figure 6(b). CMAQ produces estimates of a pollutant concentration in terms of average over a grid cell (in this case, over 10,504 12-km grid cells), while observations are taken at points. Therefore, there is a spatial misalignment between observational and numerical model data. We solve this difference in spatial resolution between the monitoring and the numerical model data by associating to each site s in the spatial domain \mathcal{S} , the CMAQ grid cell B in which s lies. This allows a comparison between the model output and the monitoring data.

To reveal the need for calibration in the numerical model output, we have produced pairwise scatterplots of the square root of the observed ozone concentration versus the square root of the CMAQ predicted concentration level of ozone during the period June 1 - September 30, 2002, of the square root of the observed ozone concentration versus the log of the CMAQ predicted $PM_{2.5}$ concentration, and analogous plots for the log of the observed concentration of particulate matter. We see that the predictions by the numerical model are biased and need to be calibrated; however, they do contain useful information to improve prediction for both pollutants. Additionally, the positive and substantial correlation between the CMAQ model output for $PM_{2.5}$ and the observed ozone concentration ($r=0.62$), and, similarly, between the CMAQ model output for ozone and the observed $PM_{2.5}$ concentration ($r=0.69$), indicate that the CMAQ output for $PM_{2.5}$ might be useful to predict ozone and conversely.

The type of calibration implied by the pairwise scatterplots mentioned above is constant across the entire spatial domain \mathcal{S} in consideration. Empirical evidence suggests instead that the error and the bias of the numerical model output might not be constant in space, rather they vary from site to site. Figure 4 presents spatial maps of the estimates of the intercept and of the regression coefficients of CMAQ ozone and CMAQ PM_{2.5} at each of the sites used to fit our bivariate downscaler model. These estimates have been obtained by regressing at each site \mathbf{s} , respectively, the normalized log of the observed PM_{2.5} concentration at \mathbf{s} across time on the corresponding normalized CMAQ model output for ozone and PM_{2.5}, both taken on the appropriate scale (square root for ozone, log for PM_{2.5}). As the plots indicate, there is spatial variability in the estimates of these coefficients with differential variability across the estimates. In particular, we see less spatial variability in the estimates obtained for the observed log PM_{2.5} concentration. Also, there is a difference between the two pollutants in the significance of the estimates of the coefficients: while for ozone, the estimates of the coefficient of CMAQ ozone are all significant, and about 60% of the estimates of the intercept are significant, in the regression for PM_{2.5}, only 21% of the intercepts are significantly different from zero, and about 88% of the estimates of the coefficient of CMAQ PM_{2.5} are significantly different from zero.

3 Modeling

Here we present the model underlying our downscaling technique. We first briefly review the univariate downscaler, presented in Berrocal et al. (2009), then we extend it to the bivariate setting. In doing so, we first present the general bivariate downscaler in a purely spatial setting, then we examine extension to accommodate data collected over time.

3.1 Univariate downscaler

Let $Y(\mathbf{s})$ be the observed data at \mathbf{s} and let $x(B)$ be the numerical model output at grid cell B . Then, $x(B)$ is interpreted as the average level of the variable under consideration over B . The spatial downscaler addresses the difference in spatial resolution between the observational data and the numerical model output, i.e. the “change of support” problem (Cressie 1993; Banerjee et al. 2004), by associating to each point \mathbf{s} the grid cell B in which \mathbf{s} lies. Then, it relates the observational data and the numerical model output as follows:

$$Y(\mathbf{s}) = \tilde{\beta}_0(\mathbf{s}) + \tilde{\beta}_1(\mathbf{s}) \cdot x(B) + \epsilon(\mathbf{s}) \quad \epsilon(\mathbf{s}) \stackrel{iid}{\sim} N(0, \tau^2) \quad (1)$$

where

$$\begin{aligned} \tilde{\beta}_0(\mathbf{s}) &= \beta_0 + \beta_0(\mathbf{s}) \\ \tilde{\beta}_1(\mathbf{s}) &= \beta_1 + \beta_1(\mathbf{s}) \end{aligned} \quad (2)$$

In (2), β_0 and β_1 denote, respectively, the overall additive and multiplicative bias of the numerical model, and $\beta_0(\mathbf{s})$ and $\beta_1(\mathbf{s})$ are local adjustments to the overall bias terms. Finally, $\epsilon(\mathbf{s})$ is a white noise process with nugget variance τ^2 . Evidently, a constant-mean downscaler is a special case of (2) obtained when the local adjustments, $\beta_0(\mathbf{s})$ and $\beta_1(\mathbf{s})$, are set to zero.

Anticipating association between intercept and slope, at the second hierarchical level of the model, we set the two spatially-varying coefficients $\beta_0(\mathbf{s})$ and $\beta_1(\mathbf{s})$ to be a bivariate correlated mean-zero Gaussian processes. Using the method of coregionalization (Wackernagel 2003; Schmidt and Gelfand 2003; Gelfand et al. 2004), we assume that there exist two latent mean-zero, unit-variance independent Gaussian processes $w_0(\mathbf{s})$ and $w_1(\mathbf{s})$ such that $\text{Cov}(w_j(\mathbf{s}), w_j(\mathbf{s}')) = \exp(-\phi_j \|\mathbf{s} - \mathbf{s}'\|)$, where, for $j = 0, 1$, ϕ_j is the spatial decay

parameter of the Gaussian process $w_j(\mathbf{s})$, $\|\mathbf{s} - \mathbf{s}'\|$ is the great circle distance between \mathbf{s} and \mathbf{s}' , and

$$\begin{pmatrix} \beta_0(\mathbf{s}) \\ \beta_1(\mathbf{s}) \end{pmatrix} = \mathbf{A} \begin{pmatrix} w_0(\mathbf{s}) \\ w_1(\mathbf{s}) \end{pmatrix} \quad (3)$$

The coregionalization matrix \mathbf{A} in (3) determines the correlation between the two spatially-varying coefficients $\beta_0(\mathbf{s})$ and $\beta_1(\mathbf{s})$ and can be assumed, without loss of generality, to be lower-triangular. Note that $\mathbf{T} = \mathbf{A}\mathbf{A}'$ is the constant local covariance matrix. The specification of the model is then completed with the priors for: the overall bias terms, β_0 and β_1 , the nugget variance τ^2 , the three non-null elements of the coregionalization matrix \mathbf{A} , and the decay parameters ϕ_0 and ϕ_1 . Other choices for obtaining dependent Gaussian processes could be entertained, including moving averages (Hoef and Barry 1998), convolution of covariance functions (Majumdar and Gelfand 2007), and particular parametric choices. In the last case, some possibilities include the cross-covariance functions introduced by Apanasovich and Genton (2010) and Gneiting et al. (2009).

3.2 The bivariate downscaler: static version

Turning to the bivariate case, we describe the model in terms of square root ozone and log $\text{PM}_{2.5}$ concentrations, recognizing that the model could be applied to any suitable pair of point referenced variables. Let $Y_1(\mathbf{s})$ and $Y_2(\mathbf{s})$ denote, respectively, the square root and the logarithm of the observed ozone and $\text{PM}_{2.5}$ concentration at \mathbf{s} , and let $x_1(B)$ and $x_2(B)$ indicate, respectively, the square root and the logarithm of the numerical model output for ozone and $\text{PM}_{2.5}$ at grid cell B . Then, $\mathbf{Y}(\mathbf{s}) = (Y_1(\mathbf{s}), Y_2(\mathbf{s}))$ is the bivariate random vector with the observed concentrations for the two pollutants at \mathbf{s} , while $\mathbf{x}(B) = (x_1(B), x_2(B))$ is the bivariate vector with the two numerical model outputs. Building upon (1), our bivariate downscaler model relates the monitoring data at \mathbf{s} and the numerical model outputs at grid cell B , with \mathbf{s} lying in B , as follows:

$$\begin{aligned} Y_1(\mathbf{s}) &= \tilde{\beta}_{10}(\mathbf{s}) + \tilde{\beta}_{11}(\mathbf{s}) x_1(B) + \tilde{\beta}_{12}(\mathbf{s}) x_2(B) + \epsilon_1(\mathbf{s}) \\ Y_2(\mathbf{s}) &= \tilde{\beta}_{20}(\mathbf{s}) + \tilde{\beta}_{21}(\mathbf{s}) x_1(B) + \tilde{\beta}_{22}(\mathbf{s}) x_2(B) + \epsilon_2(\mathbf{s}) \end{aligned} \quad (4)$$

where $\epsilon_1(\mathbf{s})$ and $\epsilon_2(\mathbf{s})$ are two independent white noise processes with nugget variances, respectively, τ_1^2 and τ_2^2 .

As in the univariate case, we adopt a random effect notation, decomposing each of the $\beta_{ij}(\mathbf{s})$, $i = 1, 2, j = 0, 1, 2$, into the sum of an overall term and a local adjustment. That is,

$$\tilde{\beta}_{ij}(\mathbf{s}) = \beta_{ij} + \beta_{ij}(\mathbf{s}). \quad (5)$$

The six-dimensional process $(\beta_{ij}(\mathbf{s}))_{i=1,2;j=0,1,2}$ is in turn modeled as a six-dimensional mean-zero correlated Gaussian process, again using the method of coregionalization. Therefore, we express each of the spatially-varying coefficients $\beta_{ij}(\mathbf{s})$ as a linear combination of mean-zero unit-variance latent independent Gaussian processes, each equipped with an exponential covariance structure. More specifically, assuming, without loss of generality, that the coregionalization matrix \mathbf{A} is lower triangular, we have that:

¹When considering large spatial domains, it is preferable to use three-dimensional Euclidean distance as an argument for the exponential covariance function rather than great-circle distance, which might yield a covariance matrix which is not positive definite. For the domain we study, the two metrics almost coincide and yield very similar results.

$$\begin{pmatrix} \beta_{10}(\mathbf{s}) \\ \beta_{11}(\mathbf{s}) \\ \beta_{12}(\mathbf{s}) \\ \beta_{20}(\mathbf{s}) \\ \beta_{21}(\mathbf{s}) \\ \beta_{22}(\mathbf{s}) \end{pmatrix} = \mathbf{A} \begin{pmatrix} w_1(\mathbf{s}) \\ w_2(\mathbf{s}) \\ w_3(\mathbf{s}) \\ w_4(\mathbf{s}) \\ w_5(\mathbf{s}) \\ w_6(\mathbf{s}) \end{pmatrix} \tag{6}$$

where, for each $k = 1, \dots, 6$, $\text{Cov}(w_k(\mathbf{s}), w_k(\mathbf{s}')) = \exp(-\phi_k \|\mathbf{s} - \mathbf{s}'\|)$, with ϕ_k decay parameter for $w_k(\mathbf{s})$. As in the univariate case, the coregionalization matrix \mathbf{A} determines, through $\mathbf{T} = \mathbf{A}\mathbf{A}'$, the correlation between the six spatially-varying coefficients $\beta_{ij}(\mathbf{s})$. Additionally, as a consequence of (4), \mathbf{A} induces a correlation structure on the bivariate random vector $\mathbf{Y}(\mathbf{s})$, i.e.

$$\Sigma_{Y(\mathbf{s}), Y(\mathbf{s}')} = [\mathbf{I}_2 \otimes (1_{x_1}(B) \ x_2(B))'] \cdot \mathbf{A} \Sigma_w \mathbf{A}' \cdot [\mathbf{I}_2 \otimes (1_{x_1}(B') \ x_2(B'))] \tag{7}$$

where B and B' are grid cells containing, respectively, \mathbf{s} and \mathbf{s}' , \mathbf{I}_2 is the identity matrix of order 2, and Σ_w is a 6×6 diagonal matrix with i -th element $\exp(-\phi_i \|\mathbf{s} - \mathbf{s}'\|)$, $i = 1, \dots, 6$.

Our bivariate downscaler model, though simple, is general and flexible. The specification of the bias of the numerical model by means of spatially-varying coefficients recognizes that calibration under the numerical model should be site-specific. Again, as in the univariate case, a constant-mean bivariate downscaler can be obtained as a special case. Moreover, permutation of the entries in the $\beta(\mathbf{s})$ vector does not affect the prior. The joint model still presents a 6×6 local covariance matrix modeled through its Cholesky decomposition.

Simplified versions of the general bivariate downscaler can be obtained by setting entries in the coregionalization matrix \mathbf{A} equal to zero.² For example, if the set of non-null entries in the \mathbf{A} matrix is given by $\{A_{11}, A_{21}, A_{22}, A_{44}, A_{64}, A_{66}\}$ and we assume that $\beta_{12} \equiv 0$ and $\beta_{21} \equiv 0$, then our bivariate downscaler reduces to two independent univariate downscalers on ozone and $\text{PM}_{2.5}$, respectively. This form of \mathbf{A} implies that the spatially-varying coefficients are correlated within pollutants but not across pollutants.

A simpler version of the bivariate downscaler that requires only three non-null entries in the coregionalization matrix \mathbf{A} but still induces a correlation among the two components $Y_1(\mathbf{s})$ and $Y_2(\mathbf{s})$ of $\mathbf{Y}(\mathbf{s})$ is the model that assumes that only the local intercept adjustments, $\beta_{10}(\mathbf{s})$ and $\beta_{20}(\mathbf{s})$, are non-null and correlated. In this case, the coregionalization matrix \mathbf{A} has only three non-null entries, $\{A_{11}, A_{41}, A_{44}\}$ and the covariance between the square root of the observed ozone concentration at \mathbf{s} and the logarithm of the observed $\text{PM}_{2.5}$ concentration at \mathbf{s}' reduces simply to:

$$\text{Cov}(Y_1(\mathbf{s}), Y_2(\mathbf{s}')) = A_{11} \cdot A_{41} \exp(-\phi_1 \|\mathbf{s} - \mathbf{s}'\|) \tag{8}$$

A simple extension to this model can be achieved by maintaining the same correlation structure for the bivariate random vector $\mathbf{Y}(\mathbf{s})$, that is, by assuming again that A_{41} is the only non-null off-diagonal element of the coregionalization matrix \mathbf{A} , but by postulating that the local adjustments to the coefficients of the two numerical model outputs, $\beta_{11}(\mathbf{s}), \beta_{12}(\mathbf{s})$,

²In principle, we could include all $21 A_{ij}$'s in the model and let the data suggest which are significant. Instead, we have chosen to do model comparison between several models, each having plausible interpretations.

$\beta_{21}(\mathbf{s})$, and $\beta_{22}(\mathbf{s})$, are non-null independent Gaussian processes with variances equal, respectively, to A_{22}^2 , A_{33}^2 , A_{55}^2 and A_{66}^2 . Then, the coregionalization matrix \mathbf{A} corresponding to this model has seven non-null entries: A_{11} , A_{22} , A_{33} , A_{44} , A_{55} , and A_{66} .

An extension of this last model that still involves fewer parameters than the general formulation for the bivariate downscaler (13 versus 21) is the model which assumes that only local adjustments of the form $\beta_{1j}(\mathbf{s})$ and $\beta_{2j}(\mathbf{s})$, for $j = 0, 1, 2$, are correlated across pollutants, while all the $\beta_{ij}(\mathbf{s})$'s specific to the same pollutant are correlated. Such a model specification corresponds to the following coregionalization matrix:

$$\mathbf{A} = \begin{pmatrix} A_{11} & 0 & 0 & 0 & 0 & 0 \\ A_{21} & A_{22} & 0 & 0 & 0 & 0 \\ A_{31} & 0 & A_{33} & 0 & 0 & 0 \\ A_{41} & 0 & 0 & A_{44} & 0 & 0 \\ 0 & A_{52} & 0 & A_{54} & A_{55} & 0 \\ 0 & 0 & A_{63} & A_{64} & 0 & A_{66} \end{pmatrix} \tag{9}$$

Other simplifications of the bivariate downscaler general model can be envisioned.

The specification of our Bayesian hierarchical bivariate downscaler model is completed with priors on the parameters. We adopt an inverse gamma with large variance for the nugget variances, τ_1^2 and τ_2^2 ; independent vague normals for the overall regression coefficients β_{ij} , $i = 1, 2, j = 0, 1, 2$; lognormals with vague standard deviations for the diagonal entries of the coregionalization matrix \mathbf{A} ; and vague normals for the off-diagonal elements of \mathbf{A} . Details on handling the decay parameters ϕ_k , $k = 1, \dots, 6$, are provided in Section 4.1.

3.3 Space-time bivariate downscaler

We now extend our general bivariate downscaler to accommodate data that has been collected over time. Let $t = 1, \dots, T$ denote the times at which we have observations and numerical model outputs, and let $Y_1(\mathbf{s}, t)$ and $Y_2(\mathbf{s}, t)$ denote, respectively, the square root and the logarithm of the observed ozone and PM_{2.5} concentration at site \mathbf{s} on day t . In an analogous way, let $x_1(B, t)$ and $x_2(B, t)$ indicate, respectively, the square root and the logarithm of the CMAQ output for ozone and PM_{2.5} on day t over the 12-km grid cell B . Then, if \mathbf{s} lies in grid cell B , the model for our bivariate downscaler becomes

$$\begin{aligned} Y_1(\mathbf{s}, t) &= \tilde{\beta}_{10}(\mathbf{s}, t) + \tilde{\beta}_{11}(\mathbf{s}, t) x_1(B, t) + \tilde{\beta}_{12}(\mathbf{s}, t) x_2(B, t) + \varepsilon_1(\mathbf{s}, t) \\ Y_2(\mathbf{s}, t) &= \tilde{\beta}_{20}(\mathbf{s}, t) + \tilde{\beta}_{21}(\mathbf{s}, t) x_1(B, t) + \tilde{\beta}_{22}(\mathbf{s}, t) x_2(B, t) + \varepsilon_2(\mathbf{s}, t) \end{aligned} \tag{10}$$

where $\varepsilon_1(\mathbf{s}, t)$ and $\varepsilon_2(\mathbf{s}, t)$ are two white noise processes that follow independently a $N(0, \tau_1^2)$ and a $N(0, \tau_2^2)$ distribution.

As in the spatial setting, for each $i = 1, 2$ and $j = 0, 1, 2$, we write:

$$\tilde{\beta}_{ij}(\mathbf{s}, t) = \beta_{ij,t} + \beta_{ij}(\mathbf{s}, t) \tag{11}$$

where, for each $t = 1, \dots, T$, $\beta_{ij}(\mathbf{s}, t)$ are correlated Gaussian processes.

Following Berrocal et al. (2009), we can model the temporal dependence in the components of (11) in several ways. For example, for the $\beta_{ij,t}$'s, we can assume that they are either

independent across time, i.e. $\beta_{ij,t} \stackrel{ind}{\sim} N(\mu_{ij}, \sigma_{ij}^2)$, or, alternatively, that they evolve dynamically in time (West and Harrison 1999), that is,

$$\beta_{ij,t} = \rho_{ij} \beta_{ij,t-1} + \eta_{ij,t}, \quad \eta_{ij,t} \stackrel{ind}{\sim} N(0, \xi_{ij}^2), \quad i=1, 2; j=0, 1, 2 \tag{12}$$

and $\beta_{ij,0} \sim N(\mu_{ij,0}, \sigma_{ij,0}^2)$.

Similar time-dependence can be imposed on the Gaussian processes $\beta_{ij}(\mathbf{s}, t)$. We can assume that the correlated Gaussian processes $\beta_{ij}(\mathbf{s}, t)$ are of the form

$$\begin{pmatrix} \beta_{10}(\mathbf{s}, t) \\ \beta_{11}(\mathbf{s}, t) \\ \beta_{12}(\mathbf{s}, t) \\ \beta_{20}(\mathbf{s}, t) \\ \beta_{21}(\mathbf{s}, t) \\ \beta_{22}(\mathbf{s}, t) \end{pmatrix} = \mathbf{A} \begin{pmatrix} w_1(\mathbf{s}, t) \\ w_2(\mathbf{s}, t) \\ w_3(\mathbf{s}, t) \\ w_4(\mathbf{s}, t) \\ w_5(\mathbf{s}, t) \\ w_6(\mathbf{s}, t) \end{pmatrix} \tag{13}$$

with \mathbf{A} still a coregionalization matrix and with the underlying Gaussian processes, $w_k(\mathbf{s}, t)$, nested within time. Then, at each $t = 1, \dots, T$, the $w_k(\mathbf{s}, t)$ are independent replicates of mean-zero unit-variance Gaussian processes with exponential correlation and decay parameters $\phi_{k,t}$ that can be either taken to be constant in time or independent across time.

In the second case, we can model the local adjustments, $\beta_{ij}(\mathbf{s}, t)$, to evolve dynamically in time. As in Gelfand et al. (2005), for each $i = 1, 2$ and $j = 0, 1, 2$, we assume that

$$\beta_{ij}(\mathbf{s}, t) = \gamma_{ij} \beta_{ij}(\mathbf{s}, t-1) + v_{ij}(\mathbf{s}, t) \tag{14}$$

where the innovations $v_{ij}(\mathbf{s}, t)$ are correlated Gaussian processes. In other words, in this second case, the coregionalization (13) is specified on the $v_{ij}(\mathbf{s}, t)$, rather than on the $\beta_{ij}(\mathbf{s}, t)$, but it still employs the underlying Gaussian processes $w_k(\mathbf{s}, t)$, which are defined as above. This dynamic model for the $\beta_{ij}(\mathbf{s}, t)$ is completed by specifying as initial conditions that $\beta_{ij}(\mathbf{s}, 0) = 0$ for $i = 1, 2$ and $j = 0, 1, 2$. In both cases we can be more general and assume that the coregionalization matrix in (13) is indexed by time, with entries independent across time. Furthermore, if t is on a continuous domain, then the $w_k(\mathbf{s}, t)$ can be independent space time processes.

The two different time dependence structures for the overall regression coefficients, $\beta_{ij,t}$, and for their local adjustments, $\beta_{ij}(\mathbf{s}, t)$, can be combined together in four ways, yielding four models that correspond to different assumptions on the way the overall and local performance and calibration of the numerical model changes over time.

3.4 The general multivariate downscaler

Extension to a multivariate downscaler given observational data and numerical model outputs for p random variables is evident. As in the bivariate case, we start with the static formulation of the model and then we extend it to the spatio-temporal setting.

Let $Y_i(\mathbf{s}), i = 1, \dots, p$ be the observed data for the i -th variable at a site \mathbf{s} in the spatial domain \mathcal{S} and let $x_i(B), i = 1, \dots, p$, be the numerical model output for the i -th variable over

grid cell B . Then, again we associate to each site \mathbf{s} the grid cell B in which \mathbf{s} lies and we relate the observational data and the numerical model output as follows:

$$Y_i(\mathbf{s}) = \tilde{\beta}_{i0}(\mathbf{s}) + \sum_{j=1}^p \tilde{\beta}_{ij}(\mathbf{s}) x_j(B) + \epsilon_i(\mathbf{s}) \quad \epsilon_i(\mathbf{s}) \stackrel{iid}{\sim} N(0, \tau_i^2) \tag{15}$$

for $i = 1, \dots, p$.

We decompose each of the $p \cdot (p + 1)$ terms $\tilde{\beta}_{ij}(\mathbf{s})$, $i = 1, \dots, p$, $j = 0, 1, \dots, p$ in the sum of an overall term and a local adjustment: $\tilde{\beta}_{is}(\mathbf{s}) = \beta_{ij} + \beta_{ij}(\mathbf{s})$, and we model the $\beta_{ij}(\mathbf{s})$'s as correlated mean-zero Gaussian processes with exponential covariance structures using again the method of coregionalization. In the general p -dimensional multivariate downscaler, the coregionalization matrix \mathbf{A} is a $p(p + 1) \times p(p + 1)$ matrix. Specifications of (15) similar to the above can be considered by simply setting to zero entries of the \mathbf{A} matrix, thus inducing simplifications on both the correlation structure of the multivariate random vector

$\mathbf{Y} = \{Y_i(\mathbf{s})\}_{i=1, \dots, p; \mathbf{s} \in \mathcal{S}}$ and on the covariance structure of the single $Y_i(\mathbf{s})$, $i = 1 \dots, p$.

To complete the specification of the model, we place standard priors on the model parameters: thus, we specify $p(p + 1)$ vague normals for the overall regression coefficients β_{ij} ; p inverse gammas with large variances for the nugget variances τ_i^2 , $i = 1 \dots, p$; log-normals with large variances for the diagonal entries A_{kk} , $k = 1, \dots, p(p + 1)$, of the coregionalization matrix \mathbf{A} ; and vague normals for the off-diagonal entries of \mathbf{A} .

If the data on the p random variables has been collected not only over space, but also over time, then the multivariate general downscaler would be extended to a spatio-temporal setting in an obvious way. Let $Y_i(\mathbf{s}, t)$, $i = 1, \dots, p$ be the observed data for the i -th random variable at site \mathbf{s} at time t , $t = 1, \dots, T$, and let $x_j(B, t)$ be the model output for the i -th variable over grid cell B at time t . Then, for $i = 1, \dots, p$, we postulate the following relationship between observational data and numerical model output:

$$Y_i(\mathbf{s}, t) = \tilde{\beta}_{i0}(\mathbf{s}, t) + \sum_{j=1}^p \tilde{\beta}_{ij}(\mathbf{s}, t) x_j(B, t) + \epsilon_i(\mathbf{s}, t) \quad \epsilon_i(\mathbf{s}, t) \stackrel{iid}{\sim} N(0, \tau_i^2) \tag{16}$$

where \mathbf{s} lies in grid cell B .

For each $i = 1, \dots, p$ and $j = 0, 1, \dots, p$, we write: $\tilde{\beta}_{ij}(\mathbf{s}, t) = \beta_{ij,t} + \beta_{ij}(\mathbf{s}, t)$, and we model the temporal structure in the $\beta_{ij,t}$ and in the $\beta_{ij}(\mathbf{s}, t)$, following what we have presented in Section 3.3 for the bivariate downscaler. Thus, we assume that the $\beta_{ij,t}$'s are either nested within time or dynamic in time, and, similarly, we postulate that the $\beta_{ij}(\mathbf{s}, t)$ are either independent replicates over time or are dynamic in time.

4 Model fitting

4.1 Priors

We have already mentioned briefly in Sections 3.2 and 3.4 the form of the prior distributions for the parameters of the bivariate and multivariate downscaler both in their static and spatio-temporal formulations. Here, we turn to the spatial decay parameters ϕ_k of the latent Gaussian processes $w_k(\mathbf{s})$ used in the coregionalization. Discrete uniform priors facilitate model fitting. Previous experience with such priors has always resulted in posterior

distributions placing highest probability on a value which is, essentially, the Restricted Maximum Likelihood (REML; Patterson and Thompson 1971; Harville 1977) estimate. Hence, in what follows, we propose to keep the spatial decay parameters fixed and equal to values that are determined by a sensitivity analysis.

More specifically, consider the simplified bivariate downscaler that uses only the two latent Gaussian processes $w_1(\mathbf{s})$ and $w_4(\mathbf{s})$ and is obtained when the coregionalization matrix \mathbf{A} has only three non-null entries, A_{11} , A_{41} , and A_{44} . The decay parameters ϕ_1 and ϕ_4 determine the way the spatial correlation decays with distance in $w_1(\mathbf{s})$ and $w_4(\mathbf{s})$, respectively, and, by consequence, in $Y_1(\mathbf{s})$ and $Y_2(\mathbf{s})$. In fact, the covariance structure of $Y_2(\mathbf{s})$ induced by such simplified version of the bivariate downscaler is given by the sum of two exponential covariance functions with decay parameters, respectively, equal to ϕ_1 and ϕ_4 , and a diagonal matrix with entries all equal to the nugget variance τ_2^2 .

To obtain a rough estimate of the magnitude of the decay parameters ϕ_1 and ϕ_4 , we have proceeded as follows. We have considered the monitoring data for the square root of ozone and the logarithm of $\text{PM}_{2.5}$ for a given day t . We have modeled the monitoring data for both variables as Gaussian processes with an unknown mean, respectively, $\mu_{Y_{1,t}}$ and $\mu_{Y_{2,t}}$, and with an exponential covariance function plus nugget effect, with decay parameters $\phi_{Y_{1,t}}$ and $\phi_{Y_{2,t}}$ and nugget variance, $\tau_{Y_{1,t}}^2$ and $\tau_{Y_{2,t}}^2$. We have estimated the parameters via Restricted Maximum Likelihood (REML; Patterson and Thompson 1971; Harville 1977) as implemented in the geoR package of R. We have repeated this operation for each day in the period June 1–September 30, 2002 and have obtained daily estimates of $\phi_{Y_{1,t}}$ and $\phi_{Y_{2,t}}$. Since histograms of those daily estimates exhibited long right tails, we have summarized those distributions by taking the median of the daily estimates, which yielded, respectively, 0.0016 and 0.00125, corresponding to ranges of, respectively, 1875 and 2400 km. We have then performed a sensitivity analysis to determine how the predictive performance of a downscaler model is affected by the magnitude of the decay parameter. Therefore, keeping the decay parameters ϕ_1 and ϕ_4 fixed and equal to the estimated medians of $\phi_{Y_{1,t}}$ and $\phi_{Y_{2,t}}$, we have fit the spatio-temporal version of the bivariate downscaler that models the overall terms $\beta_{ij,b}$ $i = 1, 2, j = 0, 1, 2$ and the local adjustments to the overall intercepts, $\beta_{i0}(\mathbf{s}, t)$, $i = 1, 2$ as independent in time (see Section 7.2 in the Supplemental material online). For each day, we have predicted levels of ozone and $\text{PM}_{2.5}$ concentration at the validation sites \mathbf{s}_0 by sampling from the posterior predictive distributions, $f(\mathbf{Y}(\mathbf{s}_0) | \{\mathbf{Y}(\mathbf{s})\}, \{\mathbf{x}(B)\})$ and subsequently backtransforming the prediction to the original scale. We have then compared the predictions with the observations by computing the predictive mean square error (PMSE), predictive mean absolute error (PMAE), empirical coverage of the 95% predictive interval and width of the 95% predictive interval. More specifically, let $Y^*(\mathbf{s}, t)$ denote the backtransformed data and predictions (recall that we used the square root on ozone concentration and the logarithm on $\text{PM}_{2.5}$ concentration), we have computed the PMSE as follows:

$$\text{PMSE} = \frac{1}{n_v} \sum_{t=1}^T \sum_{r=1}^V (\widehat{Y}^*(\mathbf{s}_r, t) - Y^*(\mathbf{s}_r, t))^2 \cdot \mathbf{I}(Y(\mathbf{s}_r, t)) \quad (17)$$

where n_v denotes the total number of observations in the validation dataset (6530 for ozone or 2559 for $\text{PM}_{2.5}$ as noted in Section 2), \mathbf{s}_r is the r -th site in the validation set containing a total of V sites, $\widehat{Y}^*(\mathbf{s}_r, t)$ denotes the posterior mean of the predictive distribution at site \mathbf{s}_r at time t on the original scale, and $\mathbf{I}(Y(\mathbf{s}_r, t))$ is equal to 1 if $Y(\mathbf{s}_r, t)$ is observed and 0 otherwise. An analogous definition holds for the PMAE, with $\widehat{Y}^*(\mathbf{s}_r, t)$ now referring to the posterior

median of the predictive distributions, again on the original scale. We have generated predictions and computed the summary statistics mentioned above four times, each time keeping the decay parameters fixed and setting them equal, respectively, to the estimated medians of $\phi_{Y_{1,t}}$ and $\phi_{Y_{2,t}}$, multiplied by 10 and 100, and divided by 10 and 100.

To compare the performance of the bivariate downscaler to that of a univariate downscaler, we have performed the same sensitivity analysis also for the univariate downscaler. Thus, keeping the decay parameters for ozone and $PM_{2.5}$ fixed and equal to the values considered for the bivariate downscaler, we have fit two spatio-temporal univariate independent downscalers for ozone and particulate matter, where the only non-null local adjustment term was the adjustment to the overall intercept, and where both the overall regression terms and the spatially varying coefficients were nested within time. Then, as in the bivariate downscaler case, we have predicted ozone and $PM_{2.5}$ concentration at the validation sites and we have assessed the performance of the predictions using the same summary statistics mentioned above.

We have carried out a similar sensitivity analysis also for ordinary kriging (Cressie 1993; Chilès and Delfiner 1999). Therefore, for both ozone and $PM_{2.5}$, for each day, using the same range of values for the decay parameters used in the experiments with the univariate and bivariate downscalers, we have kriged the observed data, on the transformed scale, to the validation sites, and we have subsequently backtransformed the predictions. Finally, we have compared the predictions with the observations using the same summary statistics that we have employed for both downscalers.

Table 1 and Table 2 report, respectively, the Predictive Mean Square Error (PMSE), Predictive Mean Absolute Error (PMAE), the empirical coverage and the width of the nominal 95% predictive interval for ozone and $PM_{2.5}$ for the two space-time downscalers and for ordinary kriging. In both tables, the results corresponding to predictions obtained when the models have been fitted using values of the decay parameters equal to the medians of the daily estimates of $\phi_{Y_{1,t}}$ and $\phi_{Y_{2,t}}$ are reported in the middle column. From both tables, it is clear that when the decay parameter is one and, especially, two orders of magnitude larger, the quality of the predictions obtained using both downscalers deteriorates noticeably in terms of PMSE and PMAE. Also, the predictive intervals are wider in this case, an indication that there is increased variability in the predictions. Misspecifying the magnitude of the decay parameters on the small side affects the quality of the predictions less noticeably. So, altogether, the sensitivity analysis has shown that both the univariate and the bivariate downscaler and ordinary kriging yield predictions that have overall best validation statistics when the decay parameters are equal to the median of the daily estimates of $\phi_{Y_{1,t}}$ and $\phi_{Y_{2,t}}$ and so, in the sequel, we fix them at these values.

4.2 Handling misalignment

As noted in Section 2, not all sites reporting ozone concentration measure $PM_{2.5}$ and vice versa. This creates a spatial misalignment in the data that we handle through the latent Gaussian processes $w_k(\mathbf{s})$ involved in the coregionalization, after having appropriately permuted and partitioned the data vector. Here, we illustrate how to handle the spatial misalignment in the static setting. Extension to the spatio-temporal case is straightforward.

Let \mathcal{S}_t be the set of n_t sites reporting measurements of ozone and/or $PM_{2.5}$ concentration on day t . We can decompose \mathcal{S}_t as the union of three disjoint sets, $\mathcal{S}_{\text{Both},t}$, $\mathcal{S}_{O,t}$, and $\mathcal{S}_{PM,t}$:

$$\mathcal{S}_t = \mathcal{S}_{\text{Both},t} \cup \mathcal{S}_{O,t} \cup \mathcal{S}_{PM,t} \quad (18)$$

where the first set includes all sites $s \in \mathcal{S}$ in which both ozone and $\text{PM}_{2.5}$ have been measured on day t , the second contains all the sites in which only ozone was measured on day t , and the last includes all sites s where only $\text{PM}_{2.5}$ has been measured on day t .

Following the decomposition in (18), we reorder and partition the random vectors $\mathbf{Y}=(Y_1(\mathbf{s}), Y_2(\mathbf{s}))_{\mathbf{s} \in \mathcal{S}_t}$ and $\mathbf{W}_k=(\mathbf{W}_k(\mathbf{s}))_{\mathbf{s} \in \mathcal{S}_t}$ in the following way: $\mathbf{Y}=(\mathbf{Y}_{\text{Both}}, \mathbf{Y}_O, \mathbf{Y}_{PM})$ and $\mathbf{w}_k=(\mathbf{w}_{k,\text{Both}}, \mathbf{w}_{k,O}, \mathbf{w}_{k,PM})$ where $\mathbf{Y}=(Y_1(\mathbf{s}), Y_2(\mathbf{s}))_{\mathbf{s} \in \mathcal{S}_{\text{Both},t}}$ and analogous definitions hold for $\mathbf{Y}_O, \mathbf{Y}_{PM}, \mathbf{w}_{k,\text{Both}}, \mathbf{w}_{k,O}$, and $\mathbf{w}_{k,PM}$.

Then it is clear that both components $Y_{1,\text{Both}}=(Y_1(\mathbf{s}))_{\mathbf{s} \in \mathcal{S}_{\text{Both},t}}$ and $Y_{2,\text{Both}}=(Y_2(\mathbf{s}))_{\mathbf{s} \in \mathcal{S}_{\text{Both},t}}$ have non-missing values and both contribute to the log-likelihood (??), while for $\mathbf{s} \in \mathcal{S}_{O,t}$ and $\mathbf{s} \in \mathcal{S}_{PM,t}$ only one component, respectively, $\mathbf{Y}_{1,O}$ and $\mathbf{Y}_{2,PM}$, has non-missing values and contributes to the log-likelihood.

The latent Gaussian processes $w_k(\mathbf{s})$ are defined on the entire spatial domain \mathcal{S} . To obtain samples from the posterior distribution of $w_k(\mathbf{s})$, within each MCMC iteration, we draw samples from the full conditionals by proceeding in the following way. If Θ denotes the collection of all model parameters (thus, $\Theta=(\tau_1^2, \tau_2^2, \{A_{kl}\}_{k,l=1,\dots,6;l \geq k}, \{\beta_{ij}\}_{i=1,2;j=0,1,2})$), for each $k=1, \dots, 6$, we:

- i. sample $\mathbf{w}_{k,\text{Both}}$ from the full conditional $\pi(\mathbf{w}_{k,\text{Both}}|\mathbf{w}_{k,O}, \mathbf{w}_{k,PM}, \mathbf{Y}_{\text{Both}}, \Theta)$;
- ii. sample $\mathbf{w}_{k,O}$ from the full conditional $\pi(\mathbf{w}_{k,O}|\mathbf{w}_{k,\text{Both}}, \mathbf{w}_{k,PM}, \mathbf{Y}_{1,O}, \Theta)$;
- iii. sample $\mathbf{w}_{k,PM}$ from the full conditional $\pi(\mathbf{w}_{k,PM}|\mathbf{w}_{k,\text{Both}}, \mathbf{w}_{k,O}, \mathbf{Y}_{2,PM}, \Theta)$;
- iv. for each $\mathbf{s} \in \mathcal{S} \setminus \mathcal{S}_t$, sample $w_k(\mathbf{s})$ from the conditional distribution $\pi(w_k(\mathbf{s})|\mathbf{w}_{k,\text{Both}}, \mathbf{w}_{k,O}, \mathbf{w}_{k,PM}, \Theta)$

By repeating steps **i-iv** within each MCMC iteration, we obtain samples from the posterior distribution of $w_k(\mathbf{s})$ for $\mathbf{s} \in \mathcal{S}$, thus solving the problem of spatial misalignment between ozone and $\text{PM}_{2.5}$.

4.3 Model selection

Using results from an analysis on the predictive performance of the univariate spatio-temporal downscaler for ozone concentration (Berrocal et al. 2009), we consider only the spatio-temporal versions of the bivariate downscaler with time-varying parameters independent across time. More specifically, we have fitted the spatio-temporal versions of the four bivariate downscalers that we have presented in Section 3.2, that is:

- i. the bivariate downscaler equivalent to two independent univariate downscalers, obtained when the coregionalization matrix \mathbf{A} has as non-null entries only $\{A_{11}, A_{21}, A_{22}, A_{44}, A_{64}, A_{66}\}$;
- ii. the bivariate downscaler with only $\{A_{11}, A_{41}, A_{44}\}$ as non-null entries in the coregionalization matrix \mathbf{A} ;
- iii. the bivariate downscaler with $\{A_{11}, A_{22}, A_{33}, A_{41}, A_{44}, A_{55}, A_{66}\}$ as non-null entries in the coregionalization matrix \mathbf{A} ; and
- iv. the bivariate downscaler with coregionalization matrix \mathbf{A} given by (9)

For each of these models, we have examined their out-of-sample predictive performance, by predicting daily concentrations of ozone and $\text{PM}_{2.5}$ at the 65 validation sites described in Section 2 and shown in Figure 1. These predictions are compared with the observed ozone

and $PM_{2.5}$ levels at the validation sites during the period June 1-September 30, 2002. Note that, for each day, predictions of ozone and $PM_{2.5}$ at a validation site s_0 are obtained by sampling from the posterior predictive distribution $f(\mathbf{Y}(s_0, t) | \{\mathbf{Y}(s, t)\}, \{\mathbf{x}(B, t)\})$ and then by subsequently transforming them back to the original scale.

To quantify the predictive performance of each bivariate downscaler we have utilized the same validation statistics used in our sensitivity analysis presented in Section 4.1. We also employ two proper scoring rules: (i) the Continuous Ranked Probability Score (CRPS; Gneiting and Raftery 2007) defined as:

$$CRPS(F, y) = \int (F(z) - 1_{\{z \geq y\}})^2 dz$$

where F is the cumulative predictive distribution function, y is the observation that materializes, 1 is the Heaviside function, that is equal to 1 if z is greater than y and 0 otherwise, and (ii) the Interval Score (Gneiting and Raftery 2007), defined for a $(1-\alpha)$ -100% predictive interval with lower bound l and upper bound u as:

$$IS(l; u; y) = \left[(u - l) + \frac{2}{\alpha} (l - y) 1_{\{y < l\}} + \frac{2}{\alpha} (u - y) 1_{\{u < y\}} \right]$$

where y denotes again the observed value.

So, for each model, we have computed the Predictive Mean Square Error (PMSE), Predictive Mean Absolute Error (PMAE), the empirical coverage of the 95% predictive interval, and the width of the 95% predictive interval to obtain an indication not only of the average bias and errors in the predictions, but also of the level of uncertainty in the predictions. Additionally, the CRPS, being a strictly proper scoring rule, provides a simultaneous assessment of the calibration and sharpness of the posterior predictive distribution whereas the Interval Score rewards predictive distributions with narrower predictive intervals while imposing a penalty for observations lying outside the predictive interval.

Finally, to determine the improvement in the predictions of ozone and $PM_{2.5}$ obtained from using our downscaling approach with respect to non-model based techniques, we have compared the downscaler with ordinary kriging and cokriging (Cressie 1993; Chilès and Delfiner 1999). For kriging, using as decay parameters for ozone and $PM_{2.5}$ concentration, respectively, 0.0016 and 0.00125, for each day, we have kriged, separately, the observed daily concentrations of ozone and $PM_{2.5}$ at the test sites to the validation sites, working on the transformed scale and then backtransforming the predictions. For cokriging, we have exploited the information contained in the copollutant and, for each day, we have cokriged the observed daily concentrations of ozone, first, and $PM_{2.5}$, afterwards, to the validation sites. Note that in performing cokriging, we had to define, not only a covariance function for both ozone and $PM_{2.5}$, but also a cross-covariance function. For both pollutants, as covariance function we used again the exponential covariance function with decay parameters, respectively, equal to 0.0016 and 0.00125. For cross-covariance function, instead, we used the cross-covariance function induced by the second version of the bivariate downscaler listed above, explicitly given in (8). The parameters for such cross-covariance function were set to be equal to the posterior means of A_{11} , A_{41} and A_{44} obtained from fitting the bivariate downscaler model in **ii**. Finally, predictive intervals and uncertainty for kriging and cokriging are based on usual formulas for the kriging and cokriging variance (Chilès and Delfiner 1999).

5 Data Analysis

We first present validation results for the different models. Since we did not find a significant improvement in the predictions of ozone and $PM_{2.5}$ obtained from versions **iii** and **iv** of the bivariate downscaler, we show only results for the versions, **i** and **ii**. To distinguish between the two, we will refer to the first as the independent downscaler, while the second will be called the bivariate downscaler. Comparing the performance of the two models allows us to establish the gain in predictive performance that can be ascribed to explicitly accounting for the correlation between the two pollutants.

From Section 4.3, for each of the models considered, we have computed the predictive mean square, the predictive mean absolute error, the empirical coverage, the width of the 95% predictive intervals, the Continuous Ranked Probability Score and the Interval Score, averaging across sites and days. In order to determine whether distance from the closest monitoring sites has an effect on the quality of the predictions, we have divided the validation sites into two groups: those that have a distance from the closest monitoring test site of less or equal than 40 kilometers, and those that are more than 40 kilometers far apart from the closest monitoring test sites. For each of the two groups, consisting, respectively, of 13 and 52 sites, we have computed the same validation statistics mentioned above.

Table 3, Table 4, and Table 5 report all these summary statistics, respectively, for all the validation sites, for the validation sites that are less or equal than 40 km distant from the closest monitoring test sites, and for the validation sites that are more than 40 km away from the closest monitoring test site.

From all tables, it is clear that the bivariate downscaler yields predictions of both ozone and $PM_{2.5}$ concentrations that are less biased than those obtained using the univariate downscaler. Additionally, the bivariate downscaler has a lower average CRPS score, indicating that its predictive distribution is sharper and better calibrated than the one corresponding to the univariate downscaler. In terms of width and empirical coverage of the 95% predictive intervals, both downscalers perform similarly and both have empirical coverage close to nominal. However, the combined assesment of both properties, which is provided by the interval score, favors again the bivariate downscaler over the univariate downscaler. A similar trend in predictive performance can be observed for kriging and cokriging: exploiting the correlation between ozone and particulate matter yield predictions that are less biased than those obtained by not accounting for it. Additionally, kriging always underestimates uncertainty in the predictions, producing predictive intervals that are too narrow and, thus, do not have the nominal coverage. This is also reflected in the CRPS and the Interval Scores: kriging always has a lower CRPS score than cokriging probably due to the smaller widths of the predictive intervals, which in turn, failing to achieve the nominal coverage are penalized in terms of Interval Score.

Comparing the bivariate downscaler with cokriging, we can see that the two methods perform equally well in terms of Interval Score; however the former outperforms the latter in terms of PMSE, PMAE, and CRPS for both ozone and $PM_{2.5}$. This follows from the fact that the bivariate downscaler not only models the correlation between ozone and $PM_{2.5}$ but also uses the information contained in the numerical model output, while cokriging does not take advantage of this additional information when predicting levels of ozone and PM at the validation sites. In addition, the bivariate downscaler predicts levels of $PM_{2.5}$ better than cokriging, likely because of the sampling frequency of $PM_{2.5}$ which renders it rather difficult to interpolate to the entire spatial domain, especially when the number of sites with observations is fairly low. Note that validation of the out-of-sample predictive performance of both downscalers and kriging/cokriging for $PM_{2.5}$ occurs mostly every three days due to

the sampling frequency of $PM_{2.5}$. In days with fewer observations from monitoring sites, this can diminish predictive gain associated with both downscalers relative to spatial interpolation techniques based solely on monitoring data.

The separation of validation sites into two groups based on their distance to the closest monitoring test site, does not reveal a difference in terms of the predictive performance of the various models. The bivariate downscaler still outperforms all the other models. However, when considering sites that are less than 40 kilometers far from the closest monitoring test site, the bivariate downscaler and cokriging predict ozone concentration equally well. This is not true for particulate matter, nor it is true when we consider validation sites that are more than 40 kilometers away from the closest monitoring test site. For this group of sites, for both ozone and $PM_{2.5}$, the predictions obtained using the bivariate downscaler are superior to the predictions obtained using any other method. Counterintuitively, for all methods, the quality of the predictions are better at validation sites that are farther from the closest monitoring test sites than at validation sites that are closer to the monitoring test sites.

Our bivariate downscaler model, as well as the univariate version, can also be used to generate predictive surfaces for both ozone and $PM_{2.5}$ concentration. Figure 5 and Figure 6 display, in panels (d), predictive surfaces of ozone and $PM_{2.5}$ concentration for June 25, 2002 obtained using our bivariate downscaler model. Each panel displays the posterior predictive mean of ozone and $PM_{2.5}$ concentration obtained by sampling from the posterior predictive distribution at each site s_0 on the CMAQ grid covering the study region. As a comparison, Figure 5 and Figure 6 present also, in panel (c), the predictive surfaces obtained by simply kriging the observed concentrations of ozone and $PM_{2.5}$ to the 10,504 CMAQ grid cells covering the region. Finally, panels (a) and (b) show respectively the observed and the CMAQ predicted concentrations for both pollutants.

As both figures show, the predictive surfaces obtained using kriging are very smooth for both pollutants, while the surfaces obtained using our bivariate downscaler show more *texture* which is expected to provide improved prediction. In particular, our bivariate downscaler method generates predictive surfaces that correct the CMAQ predictions for local bias, and display a gradient that reproduces some of the features of the CMAQ surfaces while accounting for the spatial gradient in the observations. For example, Figure 5(d) predicts lower level of ozone than CMAQ in the South-Eastern corner of map, while Figure 6(d) pushes up the level of $PM_{2.5}$ predicted by CMAQ in the same region.

In order to visually quantify the local biases of the CMAQ predicted concentrations of ozone and $PM_{2.5}$ for June 25, 2002, Figure 7 displays spatial maps of the posterior predictive means of $\beta_{10}(s)$ and $\beta_{20}(s)$. Since the bivariate downscaler models has been developed on the transformed scale, immediate conversion of these maps into calibration for the numerical model is not possible. However, they do display the spatial variability in the additive bias of the numerical model output as well as identify areas where the numerical model either underpredicts or overpredicts a pollutant concentration.

An additional feature of our downscalers is the possibility to upscale and predict the variable of interest at any spatial scale. To illustrate such capability for our bivariate downscaler, we have divided our study region into three subregions: the Mid-Atlantic (Virginia, West Virginia, Maryland, Delaware, and Pennsylvania), the Midwest (Kentucky, Ohio, Illinois, Indiana, and Missouri), and the South (all the remaining states contained in the study region). For each of these regions, we have predicted the average ozone and $PM_{2.5}$ concentration on June 25, 2002, by predicting on a grid of points using the posterior predictive distribution induced by our bivariate downscaler model and by averaging the

predictions over each of the three regions. Those predictions were then used to consider contrasts that might be of interest. Hence, Figure 8 presents plots of the posterior predictive distribution of the contrast in average ozone and $PM_{2.5}$ concentration, respectively, between the Mid-Atlantic and the Midwest, Mid-Atlantic and the South and the Midwest and the South for June 25, 2002. As the Figure shows, both the average ozone and $PM_{2.5}$ concentration are higher in the Mid-Atlantic region than in the Midwest and the South on June 25, 2002, with the South having, on average, the lowest average concentration for both pollutants on the selected day. We could further push such comparisons to look at averages or contrasts over time.

6 Summary

We have proposed a bivariate downscaler model that allows us to scale down the outputs of a numerical model from grid level to point level. The model is rather general and flexible and can easily be extended to a multivariate downscaler model for p numerical model outputs. For the bivariate case, we have examined several simplifications of the general bivariate downscaler model, each inducing a different type of correlation structure among the two components of the bivariate random vector. We have also shown how the model can be extended to accommodate data collected over both space and time. Also, as a process model, by suitable block averaging, inference can be scaled up to provide average exposure over arbitrary regions. In addition, we have demonstrated the computational advantages of scaling down, due to the relative sparsity of monitoring data compared with model output data, and the computational feasibility for our models with large space-time datasets.

In our case study, we have applied the bivariate downscaler to downscale and predict ozone and $PM_{2.5}$ concentration during the summer months of year 2002; however it is clear that our model can be used for any pair of pollutants. For example, it would be interesting to apply the bivariate downscaler model to predict wet and dry sulfate deposition, embedding the bivariate downscaling approach in the modeling context proposed by Sahu et al. (2009). Another potentially interesting analysis would be to work with O_3 and CO.

There are two extensions we are currently exploring. First, in both the univariate and the bivariate downscaler model, the numerical model output has been taken as data and we have not accounted for uncertainty in the predictions given by the numerical model. Here, potential issues involve too much smoothing of the CMAQ output, potential displacement of CMAQ grid cells, and errors in inputs that drive CMAQ. Within our downscaling approach it is possible to characterize the uncertainty in the numerical model output by adding another hierarchical level to our framework. That is, we return to the issue in the Introduction where we now require a stochastic model for the CMAQ output.

Secondly, we seek extension beyond Gaussian specifications. For instance, if we were to study extreme fields for two pollutants from both monitoring data and CMAQ, we would not use Gaussian models. Similarly, if we recorded two binary variables at each location, resulting in a 2×2 table for the location, again, Gaussian processes are not appropriate.

7 Appendix

7.1 Covariance and cross-covariance

Here we present explicit formulas for the covariance and cross-covariance between $Y_1(\mathbf{s})$ and $Y_2(\mathbf{s})$ for the different versions of the bivariate downscaler considered in Section 3.2.

For the bivariate downscaler with coregionalization matrix \mathbf{A} with non-null entries, A_{11} , A_{41} , A_{44} , the covariance between the square root of the observed ozone concentration, $Y_1(\mathbf{s})$ and

the logarithm of the observed PM_{2.5} concentration at \mathbf{s}' is given by (8). Furthermore, for this model:

$$\begin{aligned}\text{Cov}(Y_1(\mathbf{s}), Y_1(\mathbf{s}')) &= A_{11}^2 \exp(-\phi_1 \|\mathbf{s} - \mathbf{s}'\|) + \tau_1^2 \delta_{\mathbf{s}, \mathbf{s}'} \\ \text{Cov}(Y_2(\mathbf{s}), Y_2(\mathbf{s}')) &= A_{41}^2 \exp(-\phi_1 \|\mathbf{s} - \mathbf{s}'\|) + A_{44}^2 \exp(-\phi_4 \|\mathbf{s} - \mathbf{s}'\|) + \tau_1^2 \delta_{\mathbf{s}, \mathbf{s}'}\end{aligned}$$

where $\delta_{\mathbf{s}, \mathbf{s}'}$ is equal to 1 if $\mathbf{s} \equiv \mathbf{s}'$ and equal to 0 otherwise.

For the bivariate downscaler with coregionalization matrix \mathbf{A} with non-null entries, $\{A_{11}, A_{22}, A_{33}, A_{41}, A_{44}, A_{55}, A_{66}\}$, the cross-covariance between $Y_1(\mathbf{s})$ and $Y_2(\mathbf{s}')$ is given by (8), while the expressions for the within pollutant covariances are, respectively:

$$\begin{aligned}\text{Cov}(Y_1(\mathbf{s}), Y_1(\mathbf{s}')) &= A_{11}^2 \exp(-\phi_1 \|\mathbf{s} - \mathbf{s}'\|) + \sum_{k=2}^3 [A_{kk}^2 \cdot x_{k-1}(\mathbf{B}) \cdot x_{k-1}(\mathbf{B}') \cdot \exp(-\phi_k \|\mathbf{s} - \mathbf{s}'\|)] + \tau_1^2 \delta_{\mathbf{s}, \mathbf{s}'} \\ \text{Cov}(Y_2(\mathbf{s}), Y_2(\mathbf{s}')) &= A_{41}^2 \exp(-\phi_1 \|\mathbf{s} - \mathbf{s}'\|) + A_{44}^2 \exp(-\phi_4 \|\mathbf{s} - \mathbf{s}'\|) + \\ &\quad \sum_{k=5}^6 [A_{kk}^2 \cdot x_{(k-4)}(\mathbf{B}) \cdot x_{(k-4)}(\mathbf{B}') \cdot \exp(-\phi_k \|\mathbf{s} - \mathbf{s}'\|)] + \tau_2^2 \delta_{\mathbf{s}, \mathbf{s}'}\end{aligned}$$

with \mathbf{B} and \mathbf{B}' grid cell containing, respectively, \mathbf{s} and \mathbf{s}' .

Finally, for the bivariate downscaler model with coregionalization matrix \mathbf{A} given by (9), the covariance between the two components $Y_1(\mathbf{s})$ and $Y_2(\mathbf{s}')$ of the bivariate random vector $\mathbf{Y}(\mathbf{s})$ is given by:

$$\text{Cov}(Y_1(\mathbf{s}), Y_1(\mathbf{s}')) = A_{11}A_{41} \cdot \exp(-\phi_1 \|\mathbf{s} - \mathbf{s}'\|) + \sum_{k=2}^3 [A_{kk}A_{(k+3)k} \cdot x_{k-1}(\mathbf{B}) \cdot x_{k-1}(\mathbf{B}') \cdot \exp(-\phi_k \|\mathbf{s} - \mathbf{s}'\|)]$$

while the expression for the inter-pollutant covariances are given, respectively, by:

$$\begin{aligned}\text{Cov}(Y_1(\mathbf{s}), Y_1(\mathbf{s}')) &= A_{11} \exp(-\phi_1 \|\mathbf{s} - \mathbf{s}'\|) \cdot [A_{11} + A_{21}(x_1(\mathbf{B}) + x_1(\mathbf{B}')) + A_{31}(x_2(\mathbf{B}) + x_2(\mathbf{B}'))] + \\ &\quad \sum_{k=2}^3 [A_{kk}^2 \cdot x_{k-1}(\mathbf{B}) \cdot x_{k-1}(\mathbf{B}') \cdot \exp(-\phi_k \|\mathbf{s} - \mathbf{s}'\|)] + \tau_1^2 \delta_{\mathbf{s}, \mathbf{s}'}\end{aligned}$$

and

$$\begin{aligned}\text{Cov}(Y_2(\mathbf{s}), Y_2(\mathbf{s}')) &= A_{41}^2 \exp(-\phi_1 \|\mathbf{s} - \mathbf{s}'\|) + \sum_{k=2}^3 [A_{(k+3)k}^2 \cdot x_{k-1}(\mathbf{B}) \cdot x_{k-1}(\mathbf{B}') \cdot \exp(-\phi_k \|\mathbf{s} - \mathbf{s}'\|)] + \\ &\quad A_{44} \cdot \exp(-\phi_4 \|\mathbf{s} - \mathbf{s}'\|) \cdot [A_{44} + A_{54}(x_1(\mathbf{B}) + x_2(\mathbf{B}')) + A_{64}(x_2(\mathbf{B}) + x_2(\mathbf{B}'))] + \\ &\quad \sum_{k=5}^6 [A_{kk}^2 \cdot x_{k-4}(\mathbf{B}) \cdot x_{k-4}(\mathbf{B}') \cdot \exp(-\phi_k \|\mathbf{s} - \mathbf{s}'\|)] + \tau_2^2 \delta_{\mathbf{s}, \mathbf{s}'}\end{aligned}$$

References

- Apanasovich T, Genton M. Cross-covariance functions for multivariate random fields based on latent dimensions. *Biometrika*. 2010 in press.
- Banerjee, S.; Carlin, BP.; Gelfand, AE. *Hierarchical Modeling and Analysis for Spatial Data*. Chapman & Hall/CRC; Boca Raton, Fla.: 2004.
- Berrocal VJ, Gelfand AE, Holland DM. A spatio-temporal downscaler for outputs from numerical models. *Journal of Agricultural, Biological and Environmental Statistics*. 2009 in press.

- Braga ALF, Zanobetti A, Schwartz J. The lag structure between particulate air pollution and respiratory and cardiovascular deaths in ten U.S. cities. *Journal of Occupational Environmental Medicine*. 2001; 43:927–933.
- Brown PJ, Le ND, Zidek JV. Multivariate spatial interpolation and exposure to air pollutants. *Canadian Journal of Statistics*. 1994; 22:489–510.
- Byun D, Schere KL. Review of the governing equations, computational algorithms, and other components of the Models-3 Community Multiscale Air Quality (CMAQ) Modeling System. *Applied Mechanics Reviews*. 2006; 59:51–77.
- Carroll RJ, Chen EI, George TH, Li HJ, Newton H, Schmiediche H, Wang N. Ozone exposure and population density in Harris County, Texas (with discussion). *Journal of the American Statistical Association*. 1997; 92:392–415.
- Chilès, J-P.; Delfiner, P. *Geostatistics: Modeling Spatial Uncertainty*. Wiley; New York: 1999.
- Cressie, NAC. *Statistics for Spatial Data*. Wiley; New York: 1993.
- Daley, R. *Atmospheric Data Analysis*. Cambridge University Press; New York: 1991.
- Davis JM, Swall JL. An examination of the CMAQ simulations of the wet deposition of ammonium from a Bayesian perspective. *Atmospheric Environment*. 2006; 40:4562–4573.
- Dominici F, Peng RD, Bell ML, Pham L, McDermott A, Zeger SL, Samet JM. Fine particulate air pollution and hospital admission for cardiovascular and respiratory diseases. *Journal of the American Medical Association*. 2006; 295:1127–1134. [PubMed: 16522832]
- Dominici F, Samet JM, Zeger SL. Combining evidence on air pollution and daily mortality from the twenty largest U.S. cities: a hierarchical modeling strategy (with discussion). *Journal of the Royal Statistical Society Series A*. 2000; 163:263–302.
- Foley KM, Fuentes M. A statistical framework to combine multivariate spatial data and physical models for hurricane wind prediction. *Journal of Agricultural, Biological and Environmental Statistics*. 2008; 13:37–59.
- Fuentes M, Guttorp P, Challenor P. Statistical assessment of numerical models. *International Statistical Review*. 2003; 71:201–221.
- Fuentes M, Raftery AE. Model evaluation and spatial interpolation by Bayesian combination of observations with outputs from numerical models. *Biometrics*. 2005; 61:36–45. [PubMed: 15737076]
- Gelfand AE, Banerjee S, Gamerman D. Spatial process modelling for univariate and multivariate dynamic spatial data. *Environmetrics*. 2005; 16:465–479.
- Gelfand, AE.; Sahu, SK. Combining monitoring data and computer model output in assessing environmental exposure. In: O'Hagan, K.; West, M., editors. *Handbook of Applied Bayesian Analysis*. Oxford University Press; 2009. p. xxx-xxx.
- Gelfand AE, Schmidt AM, Banerjee S, Sirmans CF. Nonstationary multivariate process modeling through spatially varying coregionalization. *TEST*. 2004; 13:263–312.
- Gneiting, T.; Kleiber, W.; Schlather, M. *Matern cross-covariance functions for multivariate random fields*. University of Washington, Department of Statistics; 2009. Technical Report no. 549
- Gneiting T, Raftery AE. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*. 2007; 102:359–378.
- Gotway CA, Young LJ. Combining incompatible spatial data. *Journal of the American Statistical Association*. 2002; 97:632–648.
- Guillas S, Bao J, Choi Y, Wang Y. Statistical correction and downscaling of chemical transport model ozone forecasts over Atlanta. *Atmospheric Environment*. 2008; 42:1338–1348.
- Harville DA. Maximum likelihood approaches to variance component estimation and to related problems. *Journal of the American Statistical Association*. 1977; 72:320–338.
- Haslett J, Raftery AE. Space-time modeling with long-memory dependence: assessing Ireland's wind-power resource (with discussion). *Applied Statistics*. 1989; 38:1–50.
- Hoef JMV, Barry RD. Constructing and fitting model for cokriging and multivariable spatial prediction. *Journal of Statistical Planning and Inference*. 1998; 69:275–294.
- Jun M, Stein ML. Statistical comparison of observed and CMAQ modeled daily sulfate levels. *Atmospheric Environment*. 2004; 38:4427–4436.

- Kalnay, E. Atmospheric Modeling, Data Assimilation and Predictability. Cambridge University Press; 2003.
- Kennedy MC, O'Hagan A. Bayesian calibration of computer models (with discussion). *Journal of the Royal Statistical Society, Series B.* 2001; 63:425–464.
- Kibria BMG, Sun L, Zidek JV, Le ND. Bayesian spatial prediction of random space-time fields with application to mapping PM_{2.5} exposure. *Journal of the American Statistical Association.* 2002; 97:112–124.
- Le ND, Sun W, Zidek JV. Bayesian multivariate spatial interpolation with data missing by design. *Journal of the Royal Statistical Society, Series B.* 1997; 59:501–510.
- Liu, Z.; Le, N.; Zidek, JV. Combining measurements and physical model outputs for the spatial prediction of hourly ozone space-time fields. The University of British Columbia Department of Statistics; 2008. Technical Report no. 239
- Liu, Z.; Le, ND.; Zidek, JV. An appraisal of Bayesian Melding for physical-statistical modeling. The University of British Columbia Department of Statistics; 2007. Technical Report no. 233
- Majumdar A, Gelfand AE. Multivariate spatial modeling for geostatistical data using convolved covariance functions. *Mathematical Geology.* 2007; 39:225–245.
- McMillan N, Holland DM, Morara M, Feng J. Combining numerical model output and particulate data using Bayesian space-time modeling. *Environmetrics.* 2009 in press.
- Meiring W, Guttorp P, Sampson PD. Space-time estimation of grid cell hourly ozone levels for assessment of a deterministic model. *Environmental and Ecological Statistics.* 1998; 5:197–222.
- Paciorek, C.; Liu, Y. Limitations of remotely-sensed aerosol as a spatial proxy for fine particulate matter. Harvard University; 2008. Biostatistics Working Paper Series no. 89
- Patterson HD, Thompson R. Recovery of inter-block information when block sizes are unequal. *Biometrika.* 1971; 58:545–554.
- Poole D, Raftery AE. Inference for deterministic simulation models: the Bayesian melding approach. *Journal of the American Statistical Association.* 2000; 95:1244–1255.
- Sahu SK, Gelfand AE, Holland DM. Spatio-temporal modeling of fine particulate matter. *Journal of Agricultural, Biological and Environmental Statistics.* 2006; 11:61–86.
- Sahu SK. High resolution space-time ozone modeling for assessing trends. *Journal of the American Statistical Association.* 2007; 102:1221–1234. [PubMed: 19759840]
- Sahu SK. Fusing point and areal level space-time data with application to wet deposition. *Journal of the Royal Statistical Society, Series C.* 2009 in press.
- Sahu SK, Mardia KV. A Bayesian kriged-Kalman model for short-term forecasting of air pollution levels. *Journal of the Royal Statistical Society, Series C.* 2005; 54:223–244.
- Schmidt AM, Gelfand AE. A Bayesian coregionalization approach for multivariate pollutant data. *Journal of the Geophysical Research.* 2003; 108(D24):8783. doi:10.1029/2002JD002905.
- Schwartz J. Air pollution and hospital admissions for respiratory diseases. *Epidemiology.* 1996; 7:20–28. [PubMed: 8664396]
- Smith BJ, Cowles MK. Correlating point-referenced radon and areal uranium data arising from a common spatial process. *Applied Statistics.* 2007; 56:313–326.
- Smith RL, Davis JM, Sacks J, Speckman P, Styer P. Regression models for air pollution and daily mortality: analysis of data from Birmingham, Alabama. *Environmetrics.* 2000; 11:719–743.
- Smith RL, Kolenikov S, Cox LH. Spatio-temporal modeling of PM_{2.5} data with missing values. *Journal of Geophysical Research Atmosphere.* 2003; 108:D249004. doi:10.1029/2002JD002914.
- Swall JL, Davis JM. A Bayesian statistical approach for the evaluation of CMAQ. *Atmospheric Environment.* 2006; 40:4883–4893.
- Swall JL, Foley KM. The impact of spatial correlation and incommensurability on model evaluation. *Atmospheric Environment.* 2009; 43:1204–1217.
- Wackernagel, H. Multivariate Geostatistics: an introduction with applications. 3rd edition. Springer; berlin: 2003.
- West, M.; Harrison, J. Bayesian Forecasting and Dynamic Models. 2nd edition. Springer-Verlag; New York: 1999.

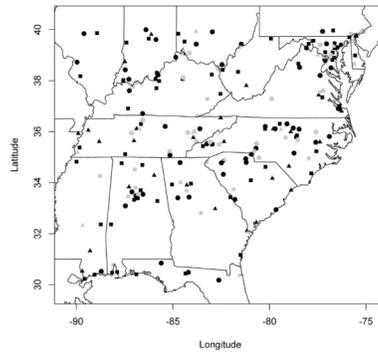


Figure 1.

Sites reporting concentration of ozone and $PM_{2.5}$ used in our case study. Sites measuring only ozone are represented with dots, sites reporting only $PM_{2.5}$ concentrations are represented with triangles, while sites measuring concentration for both pollutants are represented with squares. Black symbols are used to display sites used to fit the model, while grey symbols indicate validation sites.

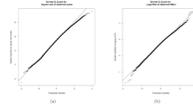


Figure 2.
Normal Q-Q plots of: (a) square root of observed ozone; (b) log of observed $PM_{2.5}$.

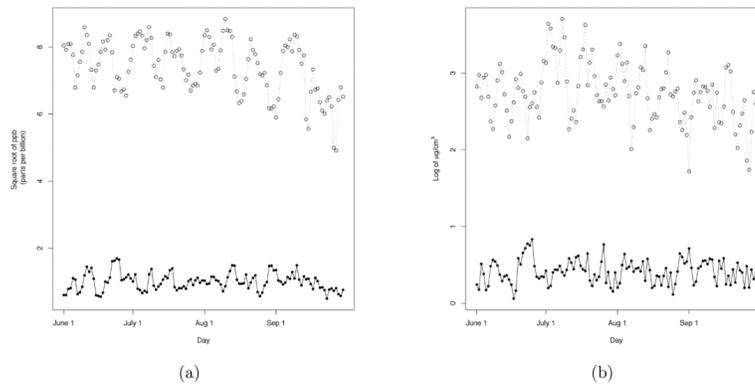


Figure 3. Time series of: (a) daily mean (open circles) and daily standard deviation (black dots) of square root of observed ozone; (b) daily mean (open circles) and daily standard deviation (black dots) of log of observed $PM_{2.5}$.

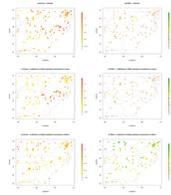


Figure 4.

Spatial maps of the estimates of the coefficients of the linear regressions of the square root of observed concentration of ozone (panels (a), (c) and (e)) and of the logarithm of the observed concentration of $PM_{2.5}$ (panels (b), (d) and (f)) on the square root of the CMAQ predicted concentration of ozone and on the log of the CMAQ predicted concentration of $PM_{2.5}$: (a)-(b) intercept term; (c)-(d) coefficient of the CMAQ model output for ozone; (e)-(f) coefficient of the CMAQ model output for $PM_{2.5}$. In all panels, the linear regression has been carried out between the normalized response and the normalized covariates.

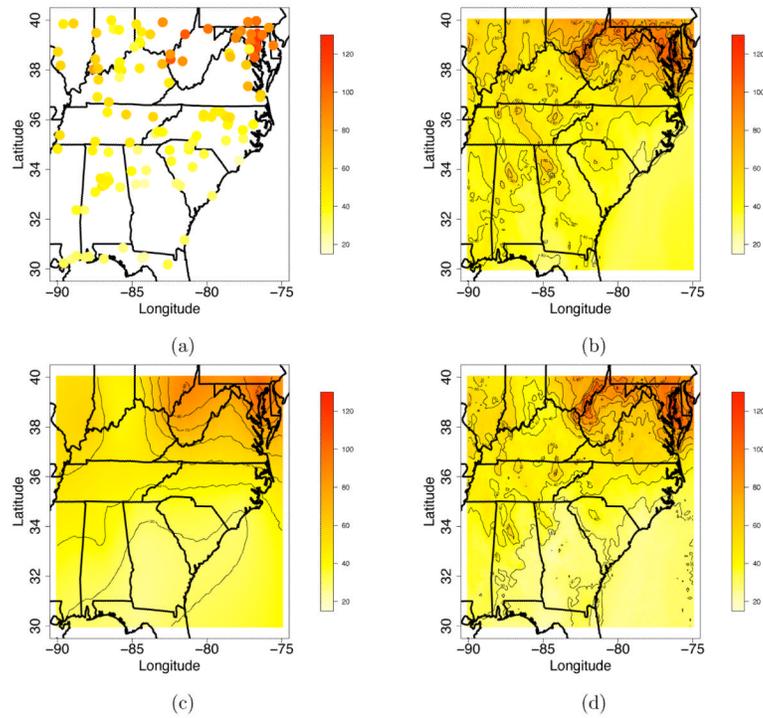


Figure 5. (a) Observed ozone on June 25, 2002; (b) Predicted ozone as obtained from CMAQ on June 25, 2002; (c) Predicted ozone as obtained via kriging for June 25, 2002; (d) Predicted ozone as obtained from the bivariate downscaler model for June 25, 2002.

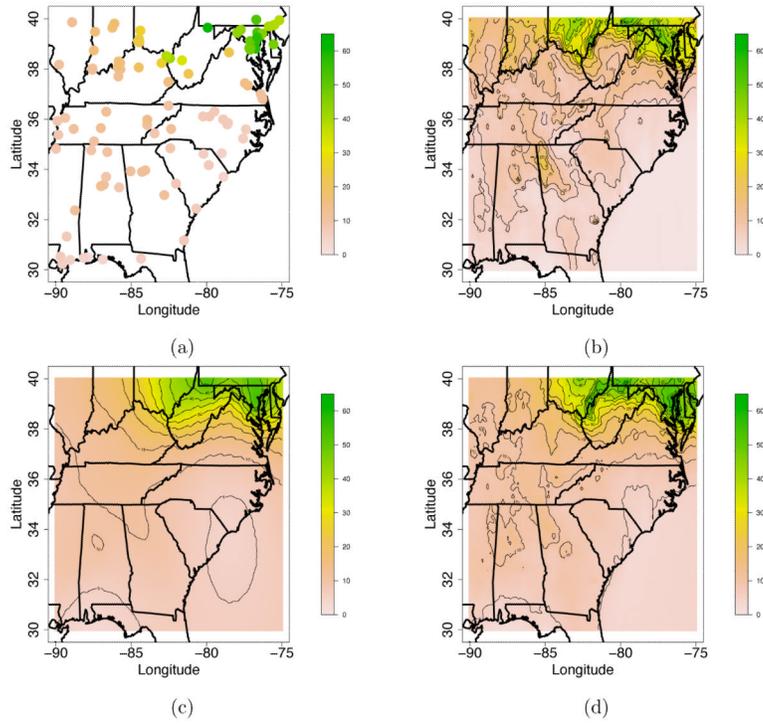


Figure 6.

(a) Observed PM_{2.5} on June 25, 2002; (b) Predicted PM_{2.5} as obtained from CMAQ on June 25, 2002; (c) Predicted PM_{2.5} as obtained via kriging for June 25, 2002; (d) Predicted PM_{2.5} as obtained from the bivariate downscaler model for June 25, 2002.

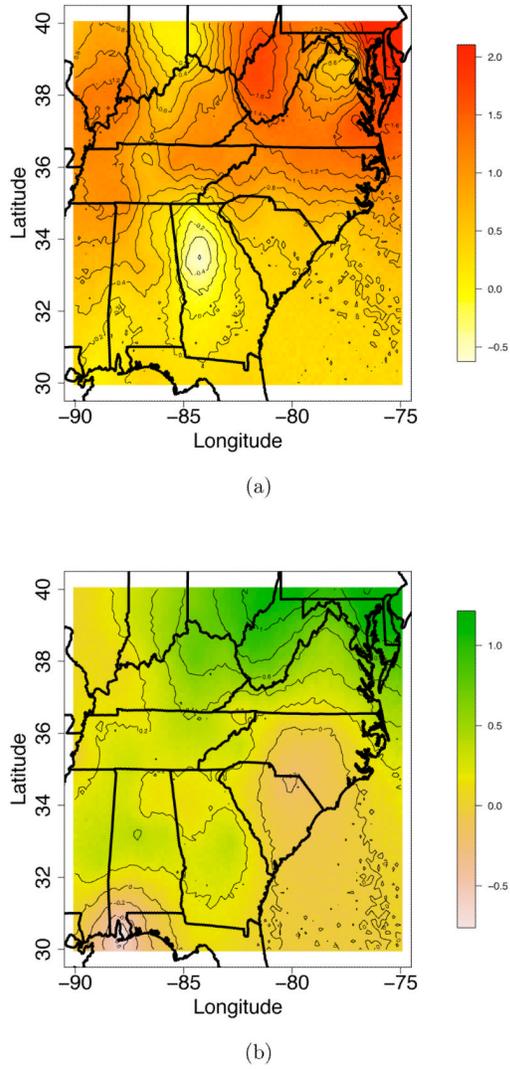


Figure 7. Posterior predictive mean of: (a) $\beta_{10}(s, t)$ for June 25, 2002; (b) $\beta_{20}(s, t)$ for June 25, 2002 as obtained from the bivariate downscaler model.

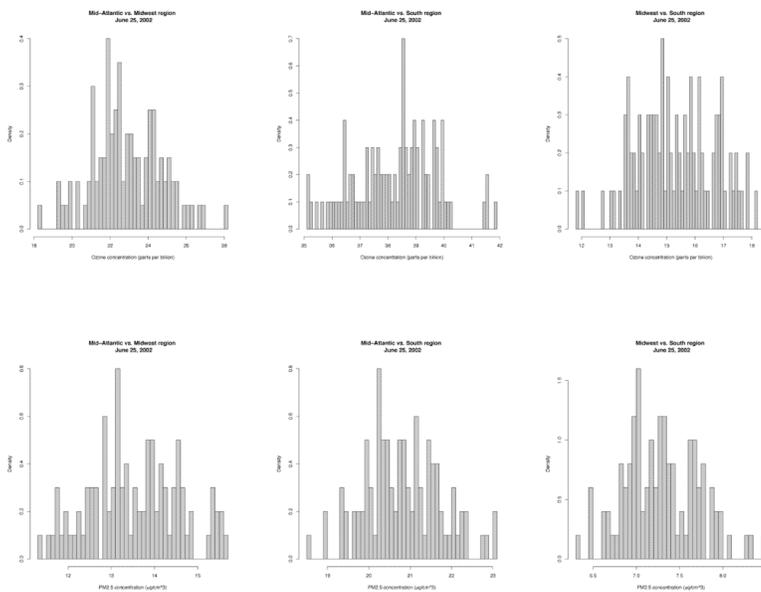


Figure 8. Posterior predictive distribution for the difference in average ozone and PM_{2.5} concentration between Mid-Atlantic and Midwest, Mid-Atlantic and South and Midwest and South on June 25, 2002, as obtained using the bivariate downscaler.

Table 1

Predictive Mean Square Error, Predictive Mean Absolute Error, empirical coverage and width of the 95% predictive interval for O_3 for different values of the decay parameter under different models

Model	Summary statistic	Decay parameter				
		0.16e-4	0.16e-3	0.16e-2	0.16e-1	0.16
Independent downscaler	PMSE	55.54	55.94	55.92	58.13	95.17
	PMAE	5.54	5.56	5.56	5.68	7.39
	Emp. cov.	93.2%	92.0%	92.9%	91.7%	92.9%
	Width	27.33	26.30	27.34	27.19	35.78
Bivariate downscaler	PMSE	53.21	51.64	52.78	55.24	91.42
	PMAE	5.45	5.36	5.40	5.54	7.27
	Emp. cov.	92.6%	92.3%	93.0%	92.7%	93.1%
	Width	26.67	26.25	26.79	27.77	34.67
Ordinary kriging	PMSE	144.58	83.28	57.16	77.19	218.03
	PMAE	9.31	6.89	5.62	6.51	11.52
	Emp. cov.	59.5%	76.6%	88.2%	92.8%	78.0%
	Width	19.10	19.73	22.26	30.26	34.69

Table 2

Predictive Mean Square Error, Predictive Mean Absolute Error, empirical coverage and width of the 95% predictive interval for $PM_{2.5}$ for different values of the decay parameter under different models

Model	Summary statistic	Decay parameter				
		0.125e-4	0.125e-3	0.125e-2	0.125e-1	0.125
Independent downscaler	PMSE	13.46	13.27	12.95	16.58	36.43
	PMAE	2.43	2.42	2.41	2.65	3.83
	Emp. cov.	93.0%	91.8%	92.7%	91.6%	92.7%
	Width	15.01	14.29	14.79	15.05	20.17
Bivariate downscaler	PMSE	11.98	11.78	11.69	14.43	32.55
	PMAE	2.33	2.31	2.29	2.49	3.55
	Emp. cov.	92.7%	93.6%	94.0%	93.2%	93.5%
	Width	14.02	14.33	14.53	14.87	18.63
Kriging	PMSE	43.89	21.99	16.19	25.65	64.17
	PMAE	4.37	3.06	2.54	3.11	5.37
	Emp. cov.	52.2%	72.5%	89.7%	94.8%	89.9%
	Width	7.32	8.89	14.06	22.61	25.73

Table 3

Predictive Mean Square Error, Predictive Mean Absolute Error, empirical coverage, width of the 95% predictive interval, Continuous Ranked Probability Score and Interval Score for both pollutants for the different models, kriging and cokriging at all the 65 validation sites

Pollutant		CMAQ	Independent downscaler	Bivariate downscaler	Kriging	Cokriging
Ozone	PMSE	122.1	55.9	52.8	57.2	53.4
	PMAE	8.5	5.6	5.4	5.6	5.4
	Emp. cov.	–	92.9%	93.0%	88.2%	94.3%
	Width	–	27.3	26.7	22.2	27.9
	CRPS	8.5	4.0	3.9	4.1	5.2
	Interval score	–	40.0	38.9	43.0	37.1
PM _{2.5}	PMSE	63.8	13.0	11.7	16.2	13.9
	PMAE	5.1	2.4	2.3	2.5	2.4
	Emp. cov.	–	92.7%	94.1%	89.7%	92.6%
	Width	–	13.7	13.3	11.9	12.8
	CRPS	5.1	1.8	1.7	1.9	2.3
	Interval score	–	19.2	17.8	21.1	19.6

Table 4

Predictive Mean Square Error, Predictive Mean Absolute Error, empirical coverage and width of the 95% predictive interval, Continuous Ranked Probability Score and Interval Score for both pollutants for the different models, kriging and cokriging at the 13 validation sites that are at no more than 40 km of distance from the closest monitoring test site

Pollutant		CMAQ	Independent downscaler	Bivariate downscaler	Kriging	Cokriging
Ozone	PMSE	124.4	57.4	54.1	57.9	54.2
	PMAE	8.5	5.6	5.4	5.6	5.5
	Emp. cov.	–	92.7%	92.7%	88.4%	94.3%
	Width	–	27.3	26.7	22.2	27.9
	CRPS	8.5	4.1	4.0	4.1	5.2
	Interval score	–	41.0	39.9	43.4	37.7
PM _{2.5}	PMSE	67.9	13.3	12.1	16.9	14.5
	PMAE	5.2	2.4	2.3	2.6	2.4
	Emp. cov.	–	92.8%	94.4%	89.9%	92.9%
	Width	–	13.8	13.4	12.1	12.9
	CRPS	5.2	1.8	1.7	1.9	2.3
	Interval score	–	19.9	18.3	22.0	18.8

Table 5

Predictive Mean Square Error, Predictive Mean Absolute Error, empirical coverage and width of the 95% predictive interval, Continuous Ranked Probability Score and Interval Score for both pollutants for the different models, kriging and cokriging at the 52 validation sites that are at more than 40 km of distance from the closest monitoring test site

Pollutant		CMAQ	Independent downscaler	Bivariate downscaler	Kriging	Cokriging
Ozone	PMSE	113.1	50.1	47.4	54.4	50.1
	PMAE	8.2	5.3	5.2	5.6	5.3
	Emp. cov.	–	93.7%	94.3%	87.5%	94.5%
	Width	–	27.1	26.5	22.2	27.8
	CRPS	8.2	3.9	3.8	4.1	5.1
	Interval score	–	36.1	34.7	41.4	35.0
PM _{2.5}	PMSE	48.5	11.6	10.0	13.5	11.4
	PMAE	4.6	2.4	2.2	2.4	2.3
	Emp. cov.	–	92.3%	92.9%	88.8%	91.4%
	Width	–	13.1	12.8	11.6	12.3
	CRPS	4.6	1.7	1.6	1.8	2.2
	Interval score	–	16.6	15.7	17.6	15.6