AperTO - Archivio Istituzionale Open Access dell'Università di Torino

Assessment and visualization of threshold exceedance probabilities in complex space-time settings: A case study of air quality in Northern Italy

This is the author's manuscript

Original Citation:

Availability:

This version is available <http://hdl.handle.net/2318/133967> since 2016-02-06T11:53:28Z

Published version:

DOI:10.1016/j.spasta.2013.04.006

Terms of use:

Open Access

Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.

(Article begins on next page)

Assessment and visualization of threshold exceedance probabilities in complex space-time settings: A case study of air quality in Northern Italy

Michela Cameletti^a, Rosaria Ignaccolo^b, Dana Sylvan^{*,c}

^a*Department of Management, Economics and Quantitative Methods,
Università degli Studi di Bergamo, Italy*

^b*Department of Economics and Statistics,
Università degli Studi di Torino, Italy*

^c*Department of Mathematics and Statistics,
Hunter College of the City University of New York, USA*

Abstract

Among the many atmospheric pollutants, fine particles are known to be particularly damaging to respiratory health, and therefore many efforts are being made worldwide to explore their spatio-temporal behavior. In this paper we focus on PM₁₀, specifically addressing the probability (or risk) that such particles will exceed potentially harmful thresholds. We combine smoothing in the time domain with spatial interpolation to model threshold exceedance probabilities and their corresponding confidence regions in a flexible framework. We then present a comprehensive study of air quality in the North-Italian region Piemonte from October 2005 through March 2006. The proposed methodology, consisting of a two-stage modeling approach followed by a block bootstrap scheme, has a myriad applications to other research fields.

Key words: block bootstrap, bootstrap percentile confidence regions, particulate matter, threshold exceedance probability.

1. Introduction

It is well known that high levels of air pollution adversely impact human and environmental health. Many research studies conducted worldwide in

*Tel. +1-212-772-5748; dsylvan@hunter.cuny.edu

the last decade have found associations between mortality rates and poor air quality, especially in urban areas. Examples include Wong et al. (2008) in Asia, Peters et al. (2000) in Central Europe, Shin et al. (2008) in Canada, Johnson and Graham (2005), McCarthy et al. (2009) and Jerrett et al. (2005) in the United States, to name only a few. A comprehensive critical review of methods addressing estimation of associations between air pollution and health is given in Dominici et al. (2003), which presented a thorough comparison of epidemiological designs and statistical approaches. In order to monitor pollution levels and prevent potential threats, many environmental agencies have issued specific standards for various pollutants. For example, the last annual report of the European Environmental Agency¹ provides information about all the major pollutants (including fine particles) and gives a thorough overview of the status and trends of air quality in Europe, based on measurements of air pollutants' concentrations as well as anthropogenic emissions. It also includes the current European air quality standards.

The negative effects that fine particles have on respiratory health have prompted many efforts in recent years to explore spatial and temporal patterns of their concentrations. In this paper, we focus on PM_{10} , specifically modeling the probability of exceeding certain potentially harmful thresholds. Many recent PM_{10} studies address the behavior of space-time trends for this pollutant (see for instance Cameletti et al. (2011), Fassò and Finazzi (2011), De Iaco et al. (2012) and the references therein); however the threshold exceedance probabilities and the corresponding uncertainties are less understood. Statistically, it is possible to model and estimate the probability that a specified value of a given pollutant will be exceeded and thus identify areas where the risk of exceeding such limit values is high. Past studies focused on trend behavior revealed that inclusion of exogenous variables may lead to better estimators and predictors of pollutant concentrations. It seems therefore natural to expect that additional information (meteorological and orographical variables, for example) might help improve daily predictions of threshold exceedance probabilities. In this paper we extend the methodology introduced in Draghicescu and Ignaccolo (2009) by adding exogenous variables in the statistical models that we use to predict and map threshold exceedance probabilities. Our preliminary study Ignaccolo et al. (2011) showed promising results. Based on daily PM_{10} measurements at 24

¹<http://www.eea.europa.eu/publications/air-quality-in-europe-2012>.

sites in the North-Italian region Piemonte, as well as exogenous variables, we obtained daily predictions of threshold exceedance probabilities over the region of interest. We detected reasonable spatial patterns for the threshold exceedance probabilities; however, the confidence regions were insufficiently informative, and there were several open questions left for future investigation.

Alternative approaches toward analyzing particulate matter involve Bayesian modeling, see for instance Sahu et al. (2006), and more recently Cameletti et al. (2012), which propose fully Bayesian space-time models for mean PM_{10} levels. Such models could be used to estimate and map threshold exceedance probabilities; however, they could not provide any estimates for the corresponding standard errors. Davison and Smith (1990) use the generalized Pareto distribution to model exceedances over high thresholds and discuss models for serial dependence in point processes with applications to wave heights and river flows. For geostatistical analyses focused on indicator kriging we refer to Guardiola-Albert and Pardo-Iguzquiza (2011) and the references therein. For an exploratory alternative based on functional data analysis to get risk maps see Ignaccolo et al. (2013).

In this paper, in addition to extending the two-stage procedure introduced in Draghicescu and Ignaccolo (2009) by adding exogenous variables, we introduce a block bootstrap scheme to provide more accurate maps, as well as bootstrap percentile confidence regions for threshold exceedance probabilities in a flexible framework. The paper is organized as follows. Section 2 describes the data used in this study. In Section 3 we present a mathematical framework to model threshold exceedance probabilities for processes with complex space-time dependencies. Our goal is to limit theoretical assumptions and give a flexible, realistic setting, appropriate for a wide variety of environmental processes. Section 4 is devoted to describing a detailed block bootstrap procedure used to correctly assess uncertainty and produce reliable confidence regions for space-time threshold exceedance probabilities. In Section 5 we show applications of the previously described methodology to our data, and conclude with the discussion in Section 6.

2. Data

We use daily PM_{10} concentrations (in $\mu\text{g}/\text{m}^3$) measured from October 1, 2005 to March 31, 2006 (182 days) by the monitoring network of the North-Italian region Piemonte at 34 sites. Measurements at 30 sites are utilized for

modeling, while data at 4 sites are kept for validation (see Figure 1). The latter have been selected to cover the most representative pollution patterns and geographic features of the region. Site A is located in Torino (altitude 260 m, population 869,312), the largest city of Piemonte, characterized by persistently high pollution levels. High and persistent PM_{10} concentrations are found at site B as well. This is located in Bra (altitude 275 m, population 29,021). Both stations C and D enjoy better air quality, with site C located in Novara (altitude 154 m, population 101,739), and station D placed in the smaller town of Serravalle Scrivia (altitude 227 m, population 6,345).

The threshold set by the European Commission (2008/50/EC directive) for PM_{10} is $50 \mu\text{g}/\text{m}^3$ and it should not be exceeded for more than 35 days a year². Fine particles in the ambient air are a serious concern in Piemonte, a region with complex orography and diverse meteorological features and atmospheric dynamics. Piemonte is situated in the western part of the Po Valley and is surrounded on three sides by the Alps. These characteristics have a strong effect on the dispersion of air pollution. For instance, weak winds leading to stagnation are commonly observed in the central part of the region, whereas the strong gusts in the mountains and valleys clear the air more rapidly. Thus, lower PM_{10} concentrations are generally observed near the Alps, while higher values are more common at lower altitudes, close to the urban areas in the plains.

In Figure 2 we show boxplots of the PM_{10} concentrations used for modeling (left panel) and validation (right panel), respectively. These distributions are generally skewed and display large variability among the monitoring sites.

We also use geographic coordinates, daily maximum mixing height, daily mean wind speed, daily emission rate of primary aerosols, and altitude. The time-varying covariates are obtained from a nested system of deterministic computer-based models implemented by the environmental agency ARPA Piemonte. For a complete description and preliminary analysis of the data we refer to Cameletti et al. (2011).

²For details see Directive 2008/50/EC of the European Parliament and of the Council of May 21, 2008 on ambient air quality and cleaner air for Europe, <http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=CELEX:32008L0050:en:NOT>; implemented in Italy on August 13, 2010.

Figure 1: Locations of 34 PM_{10} monitoring sites in Piemonte: 30 sites used for modeling labeled numerically and marked by red bullets (ordered from West to East), 4 validation sites labeled A, B, C, D and marked by blue diamonds, respectively.

3. Framework

Let $D \subset \mathbb{R}^2$ and assume that at each location $s \in D$ we observe a temporal process $X_s(t)$, for which we define the threshold exceedance probability

$$\mathbf{P}_{x_0}(t, s) = P(X_s(t) \geq x_0) \quad (1)$$

Figure 2: Distributions of PM_{10} concentrations by station (ordered from West to East); orange dots indicate means over the 182 days. Left panel shows boxplots of the 30 modeling sites, whereas the right panel displays the boxplots at the validation sites; labels correspond to those in Figure 1.

for any fixed threshold $x_0 \in \mathbb{R}$. Note that $\mathbf{P}_{x_0}(t, s)$ takes values in $[0, 1]$ and is non-increasing in x_0 . The problem of interest is to predict $\mathbf{P}_{x_0}(t, s^*)$ at location $s^* \in D$ where there are no observations for any time point t , based on observations of the process $X_s(t)$ at n time points and m spatial locations. Typically n is much larger than m , and thus one may think about the observations as long time series that are spatially correlated. The naive estimator $\hat{P}_{n,m}(x_0) = \frac{1}{mn} \sum_{i=1}^n \sum_{j=1}^m 1_{\{X_{s_j}(t_i) \geq x_0\}}$ (analogous to the empirical distribution function) is too rough and cannot capture realistically the complex spatio-temporal dynamics of the process. To overcome this challenge, as well as model threshold exceedance probabilities comprehensively at all sites, we employ the following flexible, two-step procedure.

In the *first step* we use the methodology proposed in Draghicescu and Ignaccolo (2009). For each site s , we model the temporal risks non-parametrically, by using the Nadaraya-Watson (Nadaraya (1964) and Watson (1964)) kernel estimator

$$\tilde{\mathbf{P}}_{x_0}(t, s) = \frac{\sum_{i=1}^n K\left(\frac{t_i - t}{b}\right) 1_{\{X_s(t_i) \geq x_0\}}}{\sum_{i=1}^n K\left(\frac{t_i - t}{b}\right)}, \quad (2)$$

where $t_i = i/n$, $i = 1, \dots, n$ are rescaled time points and K is the so-called kernel, a symmetric, twice continuously differentiable density function

with compact support. Rather than getting the same weight as in the naive estimator, the indicator variables are now assigned different weights, thus obtaining smooth, time-varying threshold exceedance probabilities. For an illustration see Figure 3. The tuning parameter (or bandwidth) $b = b(n)$ controls the amount of smoothing. Note that b may vary with time but should not depend on the threshold x_0 in order for the resulting estimator to be non-increasing. The threshold x_0 is considered fixed. For theoretical details and asymptotic properties of the estimator (2) we refer to Draghicescu and Ignaccolo (2009), where it is assumed that $X_s(t)$ is a time-dependent transformation of a second-order stationary Gaussian process. This is a very flexible framework, suitable for many environmental processes.

In the *second step*, we use universal kriging with exogenous variables to predict the exceedance probability field at any location $s \in D$. Since linear interpolation does not guarantee that the resulting exceedance probability estimator takes values in the interval $[0, 1]$, we first apply a 1 : 1 transformation and consider $\hat{Q}_{x_0}(t, s) = \Phi^{-1}(\hat{\mathbf{P}}_{x_0}(t, s))$ which is defined on $\mathbb{R}$, where $\Phi(\cdot)$ is the standard normal cumulative distribution function. After performing kriging on the transformed field $\hat{Q}_{x_0}(t, s)$ at site $s^* \in D$, we obtain the desired exceedance probability maps by inversion: $\hat{\mathbf{P}}_{x_0}(t, s^*) = \Phi(\hat{Q}_{x_0}(t, s^*))$. We use the following model for this space-time transformed field $\hat{Q}_{x_0}(t, s)$:

$$\hat{Q}_{x_0}(t, s) = \beta Y(t, s) + w(t, s), \quad (3)$$

where $Y(t, s)$ is a vector of exogenous variables, β denotes the corresponding coefficients, and $w(t, s)$ is a zero-mean second-order stationary spatial process for any $s \in D \subset \mathbb{R}^2$. In this step time point t is fixed and the spatial dependence of the residual process is modeled through the variogram $2\gamma(t, \|s_i - s_j\|) := \text{Var}([w(t, s_i) - w(t, s_j)]^2)$, where $\|\cdot\|$ is the Euclidean distance. After considering a number of flexible parametric variogram models, we chose to use the spherical model with nugget

$$2\gamma(t, \|s_i - s_j\|) = \rho_t + \sigma_t^2 \left(\frac{3}{2} \cdot \frac{\|s_i - s_j\|}{r_t} - \frac{1}{2} \cdot \frac{\|s_i - s_j\|^3}{r_t^3} \right), \quad (4)$$

where the nugget ρ_t captures the discontinuity (or jump) at small distances, σ_t^2 is the variance of the spatial field and r_t denotes the range. For details see Chilès and Delfiner (1999), Section 2.5.1. These parameters of the spherical variogram are estimated by weighted least squares for all time points t . The best linear unbiased predictor (BLUP) of the transformed field at location

$s^* \in D$ is obtained via universal kriging (Gaetan and Guyon, 2010, page 44) as

$$\hat{Q}_{x_0}(t, s^*) = \hat{\beta}Y(t, s^*) + \hat{w}(t, s^*). \quad (5)$$

Here $\hat{\beta}$ is the generalized least squares estimate of the trend coefficients and $\hat{w}(t, s^*) = \sum_{i=1}^m \lambda_i \hat{w}(t, s_i)$ is the simple kriging predictor, with $\hat{w}(t, s_i) = \hat{Q}_{x_0}(t, s_i) - \hat{\beta}Y(t, s_i)$. The weights λ_i , $1 \leq i \leq m$ are completely determined by the variogram parameters ρ_t , σ_t^2 , and r_t . The standard error of $\hat{Q}_{x_0}(t, s^*)$ can be also expressed in terms of the interpolation parameters λ_i . However, this standard error may not be completely accurate since the variogram parameters are estimated from the same data thus adding uncertainty. In fact, the spatial kriging predictor of the residual field is an empirical (or estimated) BLUP. Moreover, the error induced by the first step (smoothing in the time domain) is not considered. In order to account for all these uncertainty sources, in the next section we propose a block bootstrap scheme that adapts Buhlmann (2002) to this space-time setting.

4. Uncertainty assessment

As mentioned previously, the elaborate scheme we proposed for modeling spatio-temporal threshold exceedance probabilities involves a number of sources of uncertainty, not all tractable. As is the case in all statistical studies, we need to consider the tradeoff between the space-time trend and the space-time residuals. Our two-stage approach, combining nonparametric and parametric methodology, targets a wide class of processes and avoids strong theoretical assumptions that would be very hard, if not impossible to check and/or justify in practice. For instance, one common approach to addressing EBLUP's in spatial prediction problems is through conditional simulation; however, this involves specifying a parametric distribution (such as the multivariate normal) that may not be appropriate. For this reason we chose to use resampling through block bootstrap in order to assess the uncertainty in the space-time maps of the threshold exceedance probabilities. Specifically, we designed a block bootstrap scheme that takes into account the temporal evolution of the process and preserves the short-term memory of the observed PM_{10} distributions.

Many resampling schemes for dependent data have been shown to have excellent asymptotic properties and extensive applicability. Politis et al. (1999) study uniform confidence bands for marginal distributions of stationary time

series (example 3.4.7 page 89), and discuss ways to generalize bootstrap procedures to the nonstationary case (Section 4.4). A thorough review of bootstrap methods in time series is given in Haerdle et al. (2003), which notes a number of shortcomings of resampling methods under dependency. Lahiri (2003), Sections 11.5-11.7 addresses bootstrapping extremes of stationary random variables. However, our study differs in that the threshold for which we model the exceedance probability does not need to be an extreme. Canty et al. (2006) address concerns stemming from complex models where bootstrapping is applied. The paper describes initial steps for providing useful and practicable diagnostics. In our study, we have been trying to respond to the challenges posed in the aforementioned article regarding appropriate diagnostics tools. Lahiri et al. (2007) introduce a generalized plug-in rule for the empirical choice of optimal block size and prove its consistency in different bootstrap estimation scenarios.

In our study we propose the following resampling scheme based on ordered blocks. After selecting the block length l , each bootstrap iteration $r = 1, 2, \dots, R$ involves the following steps:

- sample with replacement $182/l$ blocks of consecutive days (to preserve the temporal dependence, the selected blocks of consecutive days were ordered);
- for the selected days, retrieve the corresponding PM_{10} data and covariates for the 30 sites (the PM_{10} threshold is set to $x_0 = 50 \mu\text{g}/\text{m}^3$);
- predict the bootstrap threshold exceedance probability $\hat{\mathbf{P}}_{50}^r(t, s^*)$ for the 4 validation sites using the procedure described in Section 3.

These steps are repeated $R = 1000$ times obtaining the sequence of bootstrap probabilities $\text{boot} \hat{\mathbf{P}}_{50}(t, s^*) = \{\hat{\mathbf{P}}_{50}^1(t, s^*), \dots, \hat{\mathbf{P}}_{50}^r(t, s^*), \dots, \hat{\mathbf{P}}_{50}^{1000}(t, s^*)\}$ for $t = 1, \dots, 182$ at the 4 validation sites. We denote by ${}^* \hat{\mathbf{P}}_{50}(t, s^*)$ the median of the sequence $\text{boot} \hat{\mathbf{P}}_{50}(t, s^*)$ and compute the root mean square error (RMSE) as a prediction performance index in a single validation site s^* as

$\text{RMSE}(s^*) = \sqrt{\frac{\sum_{t=1}^{182} ({}^* \hat{\mathbf{P}}_{50}(t, s^*) - \tilde{\mathbf{P}}_{50}(t, s^*))^2}{182}}$, where $\tilde{\mathbf{P}}_{50}(t, s^*)$ is the threshold exceedance probability computed on the original data. Table 1 gives these RMSE's for the 4 validation sites A, B, C and D and three choices of the block length $l = 7, 14, 26$, respectively. Since $l = 7$ minimizes the RMSE for all the validation sites, we use this block length in our applications in next section.

Site s^*	$l = 7$	$l = 14$	$l = 26$
A	0.4388	0.4624	0.4828
B	0.2995	0.3195	0.3078
C	0.3694	0.3877	0.4018
D	0.5838	0.5880	0.5974

Table 1: RMSE over 1000 bootstrap replications in the four validation sites for different bootstrap block lengths.

5. Results

5.1. Estimation

As mentioned in Section 2, we use data at $m = 30$ spatial locations (monitoring sites) and $n = 182$ time points (days). The PM_{10} threshold is set to $x_0 = 50 \mu\text{g}/\text{m}^3$. The computations were done in **R**, using the **gstat** package (Pebesma, 2004). We used the truncated Gaussian density function with support $[-1, 1]$ as the kernel function in the first step with a fixed bandwidth $b = 3$, chosen empirically.

Figure 3 shows the time series of daily PM_{10} concentrations at three sites, and the corresponding smoothed threshold exceedance probabilities $\hat{\mathbf{P}}_{50}(t, s)$ after performing the first step. It can be seen that both the patterns and the ranges are quite different. For this reason we decided to use nonparametric modeling in the time domain and used the aforementioned kernel smoothing scheme. In this way we can model these time series in a comprehensive way. Then, in the second step, spatial interpolation was performed on a 56×72 regular grid covering the whole region of Piemonte. In Figure 4 we show the implementation of the second step for four fixed time points. Left panels display the spatial trend $\hat{\beta}Y(t, s)$ in equation (5), middle panels give the fitted residual variograms, while right panels map the fitted spatial residuals $\hat{w}(t, s)$, respectively. This time we can see how the patterns vary with time as well as space. Note, for instance, that on November 19, 2005 (top panels) and March 6, 2006 (bottom panels), the spatial trend is dominant compared to the spatial residuals, which have a simple dependence structure. Here it seems that inclusion of the exogenous variables helped considerably. During December 31, 2005 and January 30, 2006 we found the reverse situation: weaker trends compared to stronger spatial variation still present in the residuals.

Figure 3: Time series of PM₁₀ concentrations at three sites (left), and corresponding estimated probability to exceed 50 $\mu\text{g}/\text{m}^3$ (right).

5.2. Uncertainty assessment

Regarding the bootstrap, we sampled with replacement $k = 26$ blocks of length $l = 7$ from the $(n - l + 1)$ possible overlapping blocks. The block length $l = 7$ was chosen based on the validation results given in Table 1. It seems that a temporal window of one week captures the meteorological and air pollution patterns well. As mentioned in Section 4, we ordered the sampled blocks in order to preserve temporal dependence. The block sampling was then repeated R times, yielding the R bootstrap samples. Bootstrap repli-

cated exceedance probability maps were obtained by performing the first and second steps (as described in Section 3) on each bootstrap sample. Based on the distribution of the R bootstrap replications, we obtained the percentile maps of the threshold exceedance probabilities, and produced bootstrap percentile confidence regions (Efron and Tibshirani, 1993, equation 13.5). We used $R = 500$ bootstrap replications. As an example, in Figure 5 we show maps of the predicted probability to exceed $50 \mu g/m^3$ of PM_{10} throughout the region under study (middle panels), together with the corresponding maps for the 10th, 25th, 75th and 90th bootstrap percentiles (yielding 50% and 80% bootstrap percentile confidence regions, respectively), during the same four days showed in Figure 4. These maps identify increased risks around the metropolitan area of Torino. The blue regions indicating low threshold exceedance probabilities could be explained by proximity to mountain ranges, the Alps in the North-West, and the Apennines in the South of Piemonte, respectively. We also note that the high-risk areas stretch towards Lombardia, the neighboring region in the East.

5.3. Other summaries

As we have shown in the previous examples, based on this two-step approach in combination with the block bootstrap scheme we can produce daily spatial maps of threshold exceedance probabilities together with corresponding confidence regions. One way of visualizing the spatio-temporal patterns could be by using animation techniques, or by selecting specific days as we did in Figure 5. It is also possible to produce other types of summaries (seasonal or yearly means for instance) that may help policy makers identify areas in need of interventions to improve air quality.

As an alternative summary, one can look at the percentage of days (out of the 182) when the probability of exceeding $50 \mu g/m^3$ is high, for instance higher than 0.75. This way, it is possible to detect regions that may not be in compliance with the European standards and where interventions are needed to improve air quality. Figure 6 displays the proportion of time (over the period under study) when $\hat{P}_{50}(t, s^*) \geq 0.75$ and 50% and 80% bootstrap confidence regions, respectively, identifying, as expected, the mountain and plateau areas as less susceptible to air pollution and the metropolitan area of Torino as potentially dangerous. These maps were based on a space-time bootstrap distribution, where each r bootstrap iteration was performed as follows:

- sample the bootstrap data with the usual procedure;

- compute the bootstrap sequence of predictions ${}_{\text{boot}}\hat{\mathbf{P}}_{50}(t, s^*)$ for all the 182 time points and all the 4032 spatial grid points;
- obtain the 4032×182 matrix of predicted threshold exceedance probabilities;
- compute (the number and) the percentage of days when the predicted probability is higher than 0.75.

These steps were repeated $R = 500$ times yielding the bootstrap distribution of the percentage of days with high exceedance probability.

6. Discussion

In this paper we describe a two-stage procedure to model spatio-temporal threshold exceedance probabilities and apply it in a comprehensive study of air quality in the North-Italian region Piemonte. We expand the methodology introduced in Draghicescu and Ignaccolo (2009) by adding exogenous variables into a flexible framework that allows for complex spatio-temporal dependencies, and is thus suitable to model wide classes of real-life processes. We also propose a resampling scheme (block bootstrap) that preserves the observed temporal short-memory, yielding confidence regions for the predicted threshold exceedance probability maps. By including geographic and meteorological information, we obtained more accurate and informative spatio-temporal maps for the predicted air pollution risks. Several sources contributed to the uncertainty associated with these maps. By using resampling we were able to realistically assess the variability associated with these predictions through bootstrap percentile confidence regions. The proposed bootstrap scheme can be adapted to produce other summaries that may be more relevant for policy makers. We show an example in Subsection 5.3. To conclude, we presented a comprehensive study of air quality in Piemonte from October 2005 through March 2006, showing high risks to exceed the legal standard limit around the main connection roads and industrial areas, as well as the biggest towns in the Po valley (toward the neighboring region Lombardia). The proposed methodology, consisting of a two-stage modeling approach combined with the block bootstrap scheme is not limited to this regional study, and could be applied in many fields, including but not limited to atmospheric sciences, demography, ecology, epidemiology, finance, medicine.

Acknowledgements

We would like to thank the editors and reviewers for reading our manuscript and for their constructive suggestions. We are very grateful to Jared Cross whose comments helped improve the presentation of this article. Ignaccolo's work was partially supported by Regione Piemonte, while Sylvan's research was funded in part by the PSC-CUNY award no. 65012-00-43.

References

- Buhlmann, P., 2002. Bootstraps for time series. *Statistical Science* 17 (1), 52–72.
- Cameletti, M., Ignaccolo, R., Bande, S., 2011. Comparing spatio-temporal models for particulate matter in Piemonte. *Environmetrics* 22, 985–996.
- Cameletti, M., Lindgren, F., Simpson, D., Rue, H., 2012. Spatio-temporal modeling of particulate matter concentration through the SPDE approach. *AStA Advances in Statistical Analysis*, 1–23.
URL <http://dx.doi.org/10.1007/s10182-012-0196-3>
- Canty, A., Davison, A., Hinkley, D., Ventura, V., 2006. Bootstrap diagnostics and remedies. *The Canadian Journal of Statistics* 34, 5–27.
- Chilès, J., Delfiner, P., 1999. *Geostatistics: Modeling Spatial Uncertainty*. Wiley.
- Davison, A. C., Smith, R., 1990. Models for exceedances over high thresholds. *Journal of the Royal Statistical Society, Series B* 52, 393–442.
- De Iaco, S., Palma, M., Posa, D., 2012. Prediction of particle pollution through spatio-temporal multivariate geostatistical analysis: spatial special issue. *AStA Advances in Statistical Analysis*, 1–18.
URL <http://dx.doi.org/10.1007/s10182-012-0199-0>
- Dominici, F., Sheppard, L., Clyde, M., 2003. Health effects of air pollution: A statistical review. *International Statistical Review/Revue Internationale de Statistique* 71, 243–276.
- Draghicescu, D., Ignaccolo, R., 2009. Modeling threshold exceedance probabilities of spatially correlated time series. *Electronic Journal of Statistics* 3, 149–164.

- Efron, B., Tibshirani, R. J., 1993. An Introduction to the Bootstrap. Chapman & Hall.
- Fassò, A., Finazzi, F., 2011. Maximum likelihood estimation of the dynamic coregionalization model with heterotopic data. *Environmetrics* 22 (6), 735–748.
- Gaetan, C., Guyon, X., 2010. Spatial Statistics and Modelling. Springer.
- Guardiola-Albert, C., Pardo-Iguzquiza, E., 2011. Compositional Bayesian indicator estimation. *Stochastic Environmental Research and Risk Assessment* 25, 835–849.
- Haerdle, W., Horowitz, J., Kreiss, J., 2003. Bootstrap methods for time series. *International Statistical Review* 71, 435–459.
- Ignaccolo, R., Ghigo, S., Bande, S., 2013. Functional zoning for air quality. *Environmental and Ecological Statistics* 20, 109–127.
- Ignaccolo, R., Sylvan, D., Cameletti, M., 2011. Modeling pollutant threshold exceedance probabilities in the presence of exogenous variables. In: *Spatial2: Spatial Data Methods for Environmental and Ecological processes*. Proceedings. pp. 1–6.
URL <http://aisberg.unibg.it/handle/10446/25243>
- Jerrett, M., Burnett, R. T., Ma, R., Pope III, C. A., Krewski, D., Newbold, K. B., Thurston, G., Shi, Y., Finklestein, N., Calle, E. E., Thun, M., 2005. Spatial analysis of air pollution and mortality in Los Angeles. *Epidemiology* 16, 727–736.
- Johnson, P. R. S., Graham, J. J., 2005. Fine particulate matter national ambient air quality standards: Public health impact on populations in the Northeastern United States. *Environmental Health Perspectives* 113, 1140–1147.
- Lahiri, S., Furukawa, K., Lee, Y., 2007. Bootstrap diagnostics and remedies. *Statistical Methodology* 4, 292–321.
- Lahiri, S. N., 2003. Resampling methods for dependent data. Springer.

- McCarthy, M. C., O'Brien, T. E., Charrier, G., Hafner, H. R., 2009. Characterization of the chronic risk and hazard of hazardous air pollutants in the United States using ambient monitoring data. *Environmental Health Perspectives* 117, 790–796.
- Nadaraya, E., 1964. On estimating regression. *Theory of Probability and Its Applications* 9, 141–142.
- Pebesma, E., 2004. Multivariable geostatistics in S: the gstat package. *Computers & Geosciences* 30, 683–691.
- Peters, A., Skorkovsky, J., Kotesovek, F., Brynda, J., Spix, C., Wichmann, H. E., Heinrich, J., 2000. Associations between mortality and air pollution in Central Europe. *Environmental Health Perspectives* 108, 283–287.
- Politis, D., Romano, J. P., Wolf, M., 1999. *Subsampling*. Springer.
- Sahu, S. K., Gelfand, A. E., Holland, D. M., 2006. Spatio-temporal modeling of fine particulate matter. *Journal of Agricultural, Biological, and Environmental Statistics* 11, 61–86.
- Shin, H. H., Stieb, D. M., Jessiman, B., Goldberg, M. S., Brion, O., Brook, J., Ramsay, T., Burnett, R., 2008. A temporal, multicity model to estimate the effects of short-term exposure to ambient air pollution on health. *Environmental Health Perspectives* 116, 1147–1153.
- Watson, G., 1964. Smooth regression analysis. *Sankhya, Series A* 26, 359–372.
- Wong, C., Vichit-Vadakan, N., Kan, H., Qian, Z., 2008. Public Health and Air Pollution in Asia (PAPA): A Multicity Study of Short-Term Effects of Air Pollution on Mortality. *Environmental Health Perspectives* 116, 1195–1202.

Figure 4: Spatial trend (left), fitted variogram of residuals (middle) and fitted spatial residuals (right); fixed time points top to bottom: November 19, 2005, December 31, 2005, January 30, 2006, March 6, 2006, respectively.

Figure 5: Maps of the predicted PM_{10} exceedance probabilities; 50% and 80% bootstrap percentile confidence regions, respectively. From top to bottom: November 19, 2005, December 31, 2005, January 30, 2006, March 6, 2006.

Figure 6: Estimated percentage of days (out of 182) with threshold exceedance probabilities higher than 0.75 (center), and corresponding 50 % and 80 % bootstrap percentile confidence regions, respectively. PM_{10} threshold is set to $50 \mu\text{g}/\text{m}^3$.