

## A spatiotemporal model for Mexico City ozone levels

Gabriel Huerta,

*University of New Mexico, Albuquerque, USA*

Bruno Sansó

*Universidad Simón Bolívar, Caracas, Venezuela, and University of California, Santa Cruz, USA*

and Jonathan R. Stroud

*University of Pennsylvania, Philadelphia, USA*

[Received January 2002. Final revision July 2003]

**Summary.** We consider hourly readings of concentrations of ozone over Mexico City and propose a model for spatial as well as temporal interpolation and prediction. The model is based on a time-varying regression of the observed readings on air temperature. Such a regression requires interpolated values of temperature at locations and times where readings are not available. These are obtained from a time-varying spatiotemporal model that is coupled to the model for the ozone readings. Two location-dependent harmonic components are added to account for the main periodicities that ozone presents during a given day and that are not explained through the covariate. The model incorporates spatial covariance structure for the observations and the parameters that define the harmonic components. Using the dynamic linear model framework, we show how to compute smoothed means and predictive values for ozone. We illustrate the methodology on data from September 1997.

**Keywords:** Bayesian inference; Exponential variogram; Kriging; Markov chain Monte Carlo methods; Spatiotemporal modelling; State space models; Tropospheric ozone

### 1. Introduction

The study of levels of tropospheric ozone is important for understanding and improving air quality in major urban areas. Environmental experts and authorities have a special interest in ozone because of its effect in diminishing health, deteriorating materials and damaging vegetation. According to environmental standards, pure air should contain less than 1% ozone and exceedingly high levels may cause eye irritation and aggravate respiratory and cardiovascular diseases.

In this paper, we focus on analysing tropospheric ozone for Mexico City, one of the most polluted cities in the world. Located at the bottom of a valley, with approximately 20 million inhabitants, Mexico City has maintained high levels of pollution during recent years mainly due to huge amounts of motor vehicle and industrial activity. In 1986, city authorities recognized the magnitude of the problem and installed a network of monitoring stations to measure ozone, carbon monoxide and hydrocarbons. The network is named Red Automática de Monitoreo Ambiental (RAMA) de la Ciudad de México. Currently, RAMA stations operate 365 days per year with short periods of interruptions for calibration of the measuring instruments. Each

*Address for correspondence:* Bruno Sansó, Department of Applied Mathematics and Statistics, Baskin School of Engineering, University of California, 1156 High Street, Santa Cruz, CA 95064, USA.  
E-mail: bruno@ams.ucsc.edu

station takes measures of pollutants automatically, second by second, and the corresponding averages per hour are reported to the public. In general, the units of the measurements are parts per billion (ppb), i.e. the concentration of the substance in a volume, where the volume is divided into 1 billion parts.

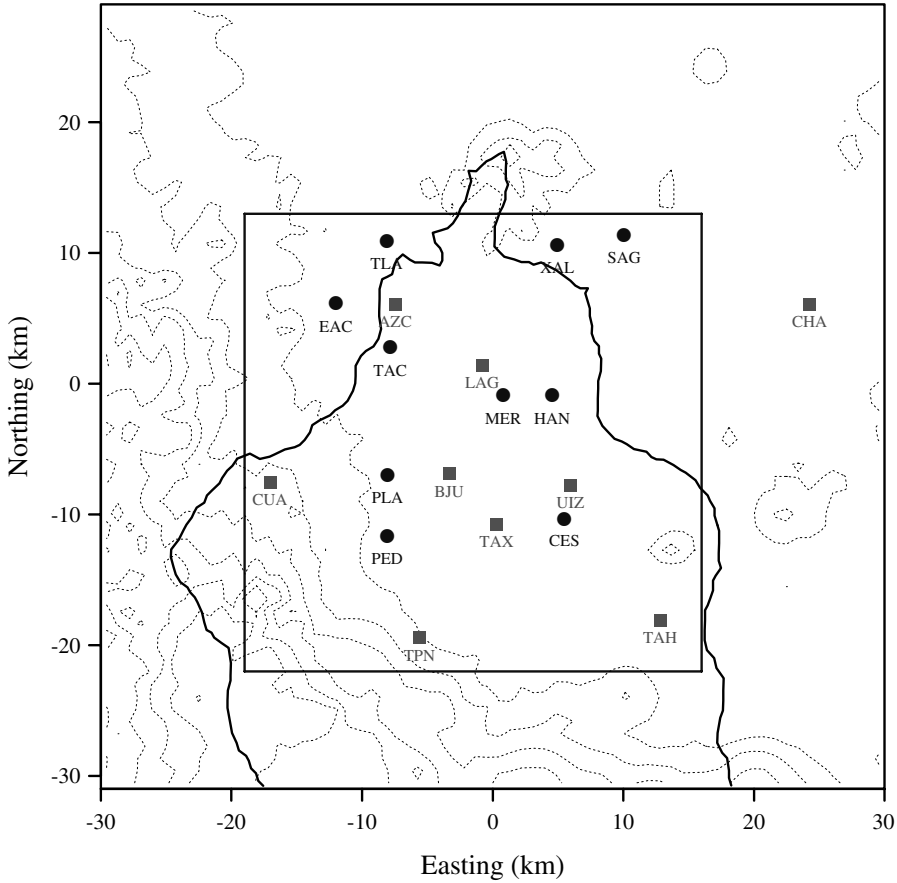
In this paper, we consider the spatiotemporal analysis of ozone time series that were obtained at some of the stations of the RAMA. As a covariate, we use temperature, which is measured only at some stations. Our goal is to propose a statistical model that forecasts temporally and interpolates spatially hourly ozone levels. Although ozone is the variable of interest, the large amounts of missing values for temperature led us to build a space–time model also for the covariate. We elaborate our models within the Bayesian paradigm by using dynamic linear models (DLMs) as in West and Harrison (1997) to account for temporal non-stationarities in the data. We use an efficient Markov chain Monte Carlo (MCMC) method to produce forecasts and spatial maps.

Previous analyses of ground level ozone data for multiple sites, modelled jointly, appear in Carroll *et al.* (1997), which used a spatially homogeneous and temporally stationary space–time model to study ozone exposure in Texas. Their model includes temperature, wind speed and wind direction as covariables. Also, Guttorp *et al.* (1994) built a space–time model for tropospheric ozone via the spatial deformation method of Sampson and Guttorp (1992) and placed it in a temporal framework by adding a stationary autoregressive process at each site. As a continuation of this work, Meiring *et al.* (1998) presented an approach to estimate hourly grid cell surface ozone concentrations in Northern California based on observations from point monitoring sites in space for an assessment of a deterministic model. This modelling approach leads to the estimation of a non-separable space–time correlation structure which is spatially non-stationary, but it involves a separate estimation of the temporal and spatial parts of the model. In contrast, there is work that considers multiple sites but modelled separately. For example, Rao *et al.* (1997) and Milanchus *et al.* (1998) considered an iterative moving average filter that decomposes ozone into a base-line, trend and a seasonal variation site by site. An extensive and critical review of different approaches of meteorological adjustment and spatiotemporal estimation of ozone is given by Thompson *et al.* (1999). Other approaches to space–time modelling appear in Stroud *et al.* (2001), Sansó and Guenni (2000), Tonellato (1997), Wikle *et al.* (1999), Berliner *et al.* (1999) and Mardia *et al.* (1998), among others. More recently, Shaddick and Wakefield (2002) proposed a multivariate spatiotemporal model of four pollutants measured daily at eight monitoring sites in London. The approach that is presented in Shaddick and Wakefield (2002) is similar to ours in that their models are developed within the dynamic linear modelling framework. However, their models are for daily, not hourly, data and assume that the effect of the potential covariates is constant in time.

The paper is organized as follows. In Section 2, we describe the data that are under study. In Section 3, we find the relevant periodicities of the ozone series by using a standard Bayesian regression tool and we discuss the analyses of the data, site by site. In Section 4, we present our space–time model for ozone, and a brief description of the Markov chain Monte Carlo (MCMC) method that is used to fit the model appears in Section 5. We leave the details of the algorithm for Appendix A. In Section 6, we present the results based on the model, and Section 7 has some conclusions and extensions of the current model.

## 2. Description of the data

We consider hourly averages of ozone in parts per billion measured during 1997 at 19 different monitoring stations scattered irregularly in Mexico City. For 10 of these 19 stations, we also



**Fig. 1.** Locations of 19 monitoring stations: ●, stations where ozone and temperature are measured; ■, stations where only ozone is measured: —, boundary of the metropolitan area of Mexico City

have hourly measurements corresponding to temperature in degrees centigrade. In Fig. 1 we show the geographical location of the 19 stations over a topographical map of Mexico City. The scale of the *x*- and *y*-axes of the map is in kilometres, with the origin in the main square of Mexico City, known as the *Zócalo*. The small box denotes the interpolation region for the space–time model of Section 5. The large box represents a region centred at *Zócalo* that contains the Mexico City metropolitan area. The dotted contour lines correspond to altitude ranging from 2200 to 2700 m.

Most of the data that we consider for our model appear in Fig. 2. The ozone time series of 11 of the stations, from September 8th to September 14th, 1997, were classified by region and plotted with the average temperature levels as a tint. The left-hand *y*-axis scale is for parts per billion and the right-hand *y*-axis is for degrees centigrade. In general, we visualize a diurnal cycle of ozone and a very high peak during the early afternoon hours, between 1 and 4 p.m. This high peak can be associated with the daily maximum temperature and the motor vehicular activity in the city during the morning and early afternoon hours. Also, there is a smaller but frequent nocturnal peak. We do not notice any obvious weekly patterns or week-end effects but there are changes from one day to the next that suggest that, even after considering daily cycles, there is a lack of stationarity in the series.

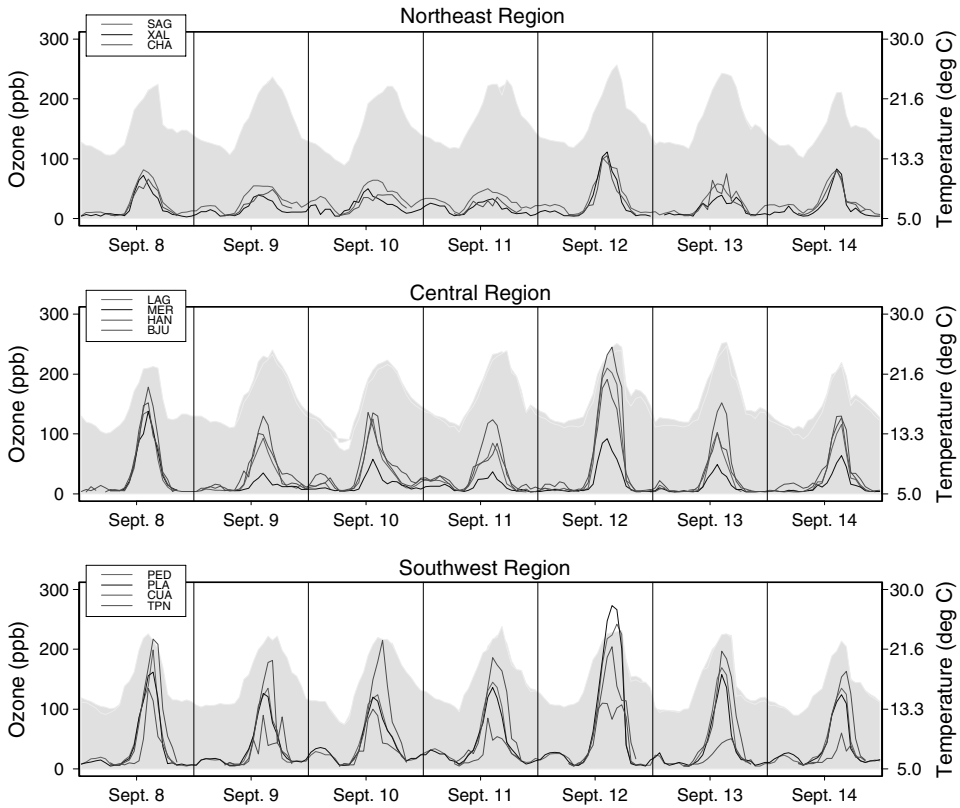


Fig. 2. Ozone and temperature time series for September 8th–14th, 1997


Fig. 2 illustrates some of the spatial variability of ozone in Mexico City. In fact, a study of the hourly means over all of 1997 for each of the 19 monitoring stations reveals that the variability of the mean level across stations is important and that the cyclical behaviour of the series differs in amplitude according to location.

As is usual for ozone measurements, the distribution of the data has an asymmetric shape that suggests the use of a transformation before building models based on the assumption of a normal distribution for the error. The two most common transformations in the literature for ozone data are the square root and the natural logarithm. Thompson *et al.* (1999) reported a summary of the transformations that have been used by different researchers who analysed ozone series. In this paper, we consider the square-root transformation for the data. This is supported by analyses of the distribution of the observed values as well as the behaviour of the residuals of the models that we propose in the following sections.

### 3. Univariate analysis

An accurate specification of the cyclical behaviour of the ozone series is a key feature for modelling such data. To make inferences on periodicities we used the *Bayesian periodogram* that was introduced by Bretthorst (1998). The Bayesian periodogram is defined as the marginal log-likelihood of the regression model

$$Y_t = a \cos(2\pi t/\lambda) + b \sin(2\pi t/\lambda) + \varepsilon_t,$$



**Fig. 3.** Bayesian periodogram of concentrations of ozone at 10 monitoring stations in Mexico City during the first 2 weeks of September 1997

with a non-informative or *reference* prior  $p(a, b, \sigma^2) \propto 1/\sigma^2$ , where  $t$  indexes time,  $\varepsilon_t \sim N(0, \sigma^2)$  and  $\lambda$  is the underlying periodicity or wavelength of the process  $Y_t$ .

Fig. 3 shows the Bayesian periodograms for the square-root ozone time series for the month of September 1997 and taken at the monitoring stations that measure temperature. The range of values for  $\lambda$  shown in Fig. 3 is between 0 and 50 h. The general pattern of all the periodograms is similar with a peak corresponding to a daily cycle with wavelengths of 24 h and a peak corresponding to a harmonic cycle with a wavelength of 12 h. Some of the stations present a smaller peak at 8 h. We think that this small peak may be related to early morning reactions between volatile organic compounds and  $\text{NO}_x$ . We also evaluated the periodograms for values of  $\lambda$  greater than 50 h and could not find any other relevant peaks. Specifically, we could not find any cycles that were associated with week-end effects.

Our initial modelling of the data consisted of relating the square-root ozone readings to the temperature by using a time series model fitted station by station. Many of the modelling strategies that were adopted for the spatiotemporal case are the product of the univariate analysis. Given the apparent non-stationarity of the data, we considered univariate dynamic regressions. This has as response variable the square root of ozone concentration  $Y_t$ , the temperature  $Z_t, t = 1, \dots, T$ , and two sine or cosine terms with periodicities  $2\pi/24$  and  $2\pi/12$  that are favoured by the Bayesian periodogram.

The model can be written as

$$\begin{aligned} Y_t &= S'_t \alpha_t + Z_t \gamma_t + \varepsilon_t, & \varepsilon_t &\sim N(0, V), \\ \alpha_t &= \alpha_{t-1} + \omega_{1t}, & \omega_{1t} &\sim N(0, \mathbf{W}_{1t}), \\ \gamma_t &= \gamma_{t-1} + \omega_{2t}, & \omega_{2t} &\sim N(0, W_{2t}) \end{aligned}$$

where  $S'_t = (\cos(\pi t/12), \sin(\pi t/12), \cos(\pi t/6), \sin(\pi t/6))$  and  $\alpha'_t = (\alpha_{1t}, \alpha_{2t}, \alpha_{3t}, \alpha_{4t})$ . The parameter  $\gamma_t$  is a scalar, and  $\varepsilon_t$ ,  $\omega_{1t}$  and  $\omega_{2t}$  are assumed independent of each other. The variance of  $\varepsilon_t$  is unknown but equal for all  $t$ , whereas the variance of  $\omega_{1t}$  and  $\omega_{2t}$  is modelled with discount factors (see West and Harrison (1997)). We disregard harmonics with higher frequencies since they do not appear relevant in the periodogram analysis of the data.

The model that was described above is a special case of a DLM as presented in West and Harrison (1997) and can be fitted by using the well-known updating and filtering equations, which are related to the Kalman filter. We fitted the model to each of the 10 stations with measurements of temperature. We used a fairly flat initial prior for the parameters at time  $t = 0$  and a global discount factor of 0.97 for the evolution covariances.

We found that this dynamic regression has a reasonable predictive performance, particularly for 24-h step ahead forecasts. Also, we notice that the estimates for  $\gamma_t$  are quite variable across time but relatively less variable across stations. In contrast, the component parameters  $\alpha_t$  are quite variable across station and time but, roughly, are such that the corresponding phase is constant across stations whereas the amplitude varies substantially.

We fitted dynamic models that did not include the two harmonic components, hoping that the periodicities in the data would be captured by the cyclical behaviour of the temperature, but the model residuals presented very high autocorrelations or severe non-normality. Similarly, we observed that the temperature coefficient is significant for most time steps and that dropping temperature resulted in a severe lack of fit. However, the addition of an 8-h harmonic, which is hinted by the periodogram, produced no significant improvements in the fit; nor did the addition of atmospheric pressure or wind as covariates.

We studied the spatial behaviour of the residuals from the station-by-station fits, plotting the empirical variograms and fitting models with geometric anisotropy that depended on the direction of the wind. None of the models that we considered produced substantial improvements over the model that consisted of a Gaussian field with isotropic exponentially decaying covariance.

#### 4. Spatiotemporal model

Let  $Y_{it}$  denote the observed square-root ozone concentration, for each station  $i$  and time  $t$ , and let  $Z_{it}$  be the temperature at time  $t$  and station  $i$ ,  $t = 1, \dots, T$ ,  $i = 1, \dots, n$ . Let  $\alpha_{it}$  denote a four-dimensional vector of seasonal coefficients for station  $i$  corresponding to a seasonal component  $S'_t$  consisting of sine and cosine terms as introduced in the previous section.

The general space–time model that we propose for ozone data is given by

$$Y_{it} = \beta_t^y + S'_t \alpha_{it} + Z_{it} \gamma_t + \varepsilon_{it}$$

where the errors  $\varepsilon_{it}$  are assumed spatially correlated with a Gaussian distribution. Here  $\beta_t^y$  defines a canonical spatial trend and  $Z_{it} \gamma_t$  is the effect of the covariate on  $Y_{it}$ . Note that the coefficient that is related to the covariate is assumed to be equal for all stations, whereas each station has its own set of seasonal parameters. This is justified by the results that were obtained with the univariate time series analyses. Clearly, this general space–time model has a very large number of parameters, since at each time  $t$  there are four parameters for each station, a common parameter for the covariate, plus the parameters that define the spatial correlation.

A substantial reduction in the number of parameters is achieved by assuming that the amplitudes of each cyclical component are different, but the phases are very similar between stations and almost constant in time. This assumption is supported by the univariate models that were fitted station by station.

We can thus consider a modification of the model given by

$$Y_{it} = \beta_t^y + S'_t(\mathbf{a})\alpha_{it} + Z_{it}\gamma_t + \varepsilon_{it}^y.$$

$\alpha_{it} = (\alpha_{1it}, \alpha_{2it})$  is now a vector of two dimensions and  $S'_t(\mathbf{a}) = (S_{1t}(a_1), S_{2t}(a_2))$  where

$$S_{jt}(a_j) = \cos(\pi jt/12) + a_j \sin(\pi jt/12),$$

$j = 1, 2$ , defines the cyclical component in the data and  $\mathbf{a} = (a_1, a_2)'$ .

Thus  $\alpha_{jit}^2(1 + a_j^2)$  is the amplitude of the  $j$ th periodicity of the  $i$ th station at time  $t$ , and  $\tan^{-1}(a_j)$  is its phase, which is independent of  $i$  and  $t$ .

In vector form, our modified space–time model can be written as

$$\mathbf{Y}_t = \mathbf{1}_n \beta_t^y + S_{1t}(a_1)\alpha_{1t} + S_{2t}(a_2)\alpha_{2t} + \mathbf{Z}_t \gamma_t + \boldsymbol{\varepsilon}_t^y$$

where  $\mathbf{Y}'_t = (Y_{1t}, \dots, Y_{nt})$ ,  $\mathbf{1}'_n = (1, \dots, 1)$ ,  $\boldsymbol{\alpha}'_{it} = (\alpha_{i1t}, \alpha_{i2t}, \dots, \alpha_{int})$ ,  $i = 1, 2$ , and  $\boldsymbol{\varepsilon}_t^y = (\varepsilon_{1t}^y, \varepsilon_{2t}^y, \dots, \varepsilon_{nt}^y)$ .

Additionally, we assume that  $\boldsymbol{\varepsilon}_t^y \sim N(0, \sigma_y^2 \mathbf{V}^y)$  with  $\mathbf{V}^y = \exp(-\mathbf{D}/\lambda_y)$ ,  $\mathbf{D}$  is the matrix of distances in kilometres between stations,  $\exp(\mathbf{D})$  denotes the exponentiation of each of its components and  $\lambda_y$  is a positive range parameter, i.e. we assume that the errors are normally distributed with an exponential covariance function.

Furthermore, the parameter evolutions of the model are defined as

$$\begin{aligned} \beta_t^y &= \beta_{t-1}^y + \omega_t^y, & \omega_t^y &\sim N(0, \sigma_y^2 \tau_y^2), \\ \boldsymbol{\alpha}_{it} &= \boldsymbol{\alpha}_{it-1} + \boldsymbol{\omega}_{it}^{\alpha_i}, & \boldsymbol{\omega}_{it}^{\alpha_i} &\sim N(0, \sigma_y^2 \tau_i^2 \mathbf{W}^{\alpha_i}), \\ \gamma_t &= \gamma_{t-1} + \omega_t^\gamma, & \omega_t^\gamma &\sim N(0, \sigma_y^2 \tau_\gamma^2) \end{aligned}$$

where  $\tau_y^2$ ,  $\tau_\gamma^2$  and  $\tau_i^2$ ,  $i = 1, 2$ , are fixed constants;  $\mathbf{W}^{\alpha_i} = \exp(-\mathbf{D}/\lambda_i)$  defines an exponential covariance function depending on the range parameter  $\lambda_i$ ,  $i = 1, 2$ . Therefore, we are incorporating spatial dependence on the parameters that determine the seasonal terms.

The above model assumes that the observations for temperature are known at each time and station so it is a model for  $\mathbf{Y}_t$  conditional on  $\mathbf{Z}_t$ . Values of  $Z_{it}$  at arbitrary locations and times are needed to interpolate the values of  $Y_{it}$  both temporally and spatially and they may not be available. Then, we also formulate a spatiotemporal model on  $Z_{it}$  defined by

$$Z_{it} = \beta_t^z + h_i \eta_t + (1, 0, 1, 0)\boldsymbol{\delta}_t + \varepsilon_{it}^z,$$

with  $\beta_t^z$  a canonical spatial trend on  $Z_{it}$ ,  $h_i$  the height of station  $i$ ,  $\eta_t$  the altitude coefficient at time  $t$ ,  $\boldsymbol{\delta}_t = (\delta_{1t}, \delta_{2t}, \delta_{3t}, \delta_{4t})$  and  $\varepsilon_{it}^z$  is the error.

In vector form, the model is

$$\mathbf{Z}_t = \mathbf{1}_n \beta_t^z + \mathbf{H}' \eta_t + \mathbf{E}' \boldsymbol{\delta}_t + \boldsymbol{\varepsilon}_t^z$$

where  $\mathbf{Z}'_t = (Z_{1t}, Z_{2t}, \dots, Z_{nt})$ ,  $\mathbf{1}'_n = (1, 1, \dots, 1)$ ,  $\mathbf{H} = (h_1, h_2, \dots, h_n)$ ,  $\mathbf{E}'$  is an  $(n \times 4)$ -dimensional matrix where all the rows are equal to the vector  $(1, 0, 1, 0)$  and  $\boldsymbol{\varepsilon}_t^z = (\varepsilon_{1t}^z, \dots, \varepsilon_{nt}^z)$  is the vector of errors with  $\boldsymbol{\varepsilon}_t^z \sim N(0, \sigma_z^2 \mathbf{V}^z)$  where  $\mathbf{V}^z = \exp(-\mathbf{D}/\lambda_z)$ .

A univariate time series analysis of the temperatures reveals that there are time changing differences in mean values due to the height of the station. In contrast the seasonality is fairly homogeneous in space. Therefore, we are considering a model for  $\mathbf{Z}_t$  with the same seasonal effect for all locations.

Additionally, the parameter evolutions are defined as

$$\begin{aligned} \beta_t^z &= \beta_{t-1}^z + \omega_t^z, & \omega_t^z &\sim N(0, \sigma_z^2 \tau_z^2), \\ \eta_t &= \eta_{t-1} + \omega_t^\eta, & \omega_t^\eta &\sim N(0, \sigma_z^2 \tau_\eta^2), \\ \delta_t &= G \delta_{t-1} + \omega_t^\delta, & \omega_t^\delta &\sim N(0, \sigma_z^2 \tau_\delta^2 \mathbf{I}) \end{aligned}$$

where  $\tau_z^2$ ,  $\tau_\eta^2$  and  $\tau_\delta^2$  are fixed constants,  $G$  is a  $4 \times 4$  block diagonal matrix with blocks of the form

$$G_j = \begin{pmatrix} \cos(\pi j/12) & \sin(\pi j/12) \\ -\sin(\pi j/12) & \cos(\pi j/12) \end{pmatrix}, \quad j = 1, 2,$$

which defines the seasonal component on  $\mathbf{Z}_t$  and  $\mathbf{I}$  is a  $4 \times 4$  identity matrix.

The model is completed by specifying a prior distribution on the range parameters of the covariance matrix, the scale parameters at the observation level and the phase parameters  $\mathbf{a}$ . A particular prior specification is presented in Section 6.

In summary, our new spatiotemporal model for ozone is defined by the two equations

$$\begin{aligned} \mathbf{Y}_t &= \mathbf{1}_n \beta_t^y + S_{1t}(a_1) \alpha_{1t} + S_{2t}(a_2) \alpha_{2t} + \mathbf{Z}_t \gamma_t + \varepsilon_t^y, \\ \mathbf{Z}_t &= \mathbf{1}_n \beta_t^z + \mathbf{H}' \eta_t + \mathbf{E}' \delta_t + \varepsilon_t^z. \end{aligned}$$

This model is in the DLM or state space form notation of West and Harrison (1997). Thus, conditional on  $\mathbf{a}$  and all the hyperparameters that define the covariance structure at both observational and evolution levels, the filtering and recurrence equations of the DLM produce predictive values and retrospective inferences for the observed data. Formal Bayesian inference on some of the hyperparameters, e.g. the range parameters  $\lambda_y$  and  $\lambda_z$ , leads to the *forward filtering–backward sampling* algorithm that we describe now.

### 5. Markov chain Monte Carlo algorithm and spatial interpolation

Under the space–time model that has just been presented, posterior, predictive and interpolation analyses are available via MCMC methods. The structure of relevant conditional distributions is briefly outlined here. Further details appear in Appendix A. First we give some notation. Define  $\theta_t^y = (\beta_t^y, \alpha_{1t}, \alpha_{2t}, \gamma_t)$  as the ozone state vector and  $\theta_t^z = (\beta_t^z, \eta_t, \delta_t)$  as the temperature state vector at time  $t$ , and set  $\theta^y = (\theta_0^y, \dots, \theta_T^y)$ ,  $\theta^z = (\theta_0^z, \dots, \theta_T^z)$ .  $\mathbf{Y} = (Y_1, \dots, Y_T)$  and  $\mathbf{Z} = (Z_1, \dots, Z_T)$ . Let the superscripts ‘o’ and ‘u’ denote the observed and unobserved data respectively. Posterior inferences are then based on summarizing the joint posterior distribution:

$$p(\mathbf{Y}^u, \theta^y, \mathbf{a}, \sigma_y^2, \lambda_y, \mathbf{Z}^u, \theta^z, \sigma_z^2, \lambda_z | \mathbf{Y}^o, \mathbf{Z}^o).$$

Our MCMC algorithm is a block sampling scheme in which we iteratively simulate each component of the joint posterior distribution. Considering blocks of parameters is a key issue to obtain convergence and computational efficiency. Related work for Markov random-field models appears in Rue (2001) and Knorr-Held and Rue (2002). The full conditional distributions are given below. The prior distributions for the unknown parameters are given in Section 6.

#### 5.1. Sampling the state vectors, error variances and range parameters

Using the state space equations for  $\mathbf{Y}_t$  and  $\mathbf{Z}_t$ , we sample  $(\theta^y, \sigma_y^2, \lambda_y)$  and  $(\theta^z, \sigma_z^2, \lambda_z)$  as blocks. Dropping the  $y$ - and  $z$ -subindices and the conditioning on all other parameters and observations, we note that



$$p(\boldsymbol{\theta}, \sigma^2, \lambda) = p(\boldsymbol{\theta}|\sigma^2, \lambda) p(\sigma^2|\lambda) q(\lambda).$$

We can obtain samples of  $p(\boldsymbol{\theta}|\sigma^2, \lambda)$  by using forward filtering–backward sampling, i.e. we apply the filtering recurrences for the DLMs defined for the response and the covariate and then sample each element of  $\boldsymbol{\theta}$  recursively from  $t = T, \dots, 1$  as in Carter and Kohn (1994) or Frühwirth-Schnatter (1994).  $p(\sigma^2|\lambda)$  corresponds to an inverse gamma distribution and  $p(\lambda)$  can be easily evaluated for each  $\lambda$ . The details appear in Appendix A. We then propose the following Metropolis–Hastings scheme.

- (a) Choose a jumping distribution for  $\lambda$ , say,  $q(\lambda^*|\lambda)$ .
- (b) Sample  $(\boldsymbol{\theta}^*, \sigma^{2*}, \lambda^*) \sim p(\boldsymbol{\theta}|\sigma^{2*}, \lambda^*) p(\sigma^{2*}|\lambda^*) q(\lambda^*|\lambda)$ .
- (c) Let

$$\alpha = \min \left\{ 1, \frac{p(\lambda^*) q(\lambda|\lambda^*)}{p(\lambda) q(\lambda^*|\lambda)} \right\}.$$

- (d) Accept  $\boldsymbol{\theta}^*, \sigma^{2*}$  and  $\lambda^*$  with probability  $\alpha$ .

Note that the acceptance ratio  $\alpha$  does not depend on  $\boldsymbol{\theta}^*$  and  $\sigma^{2*}$ , so we can avoid the backward sampling of  $\boldsymbol{\theta}$  if  $\lambda^*$  is not accepted. This method produces a substantial reduction in the computational cost of the MCMC algorithm. Our implementation of the algorithm considers a transition density  $q(\lambda^*|\lambda)$  defined on the  $\log(\lambda)$  scale that follows a normal distribution centred on the last sampled value.

**5.2. Conditional for missing values**

Since we are assuming that the error terms  $\varepsilon_t^y$  and  $\varepsilon_t^z$  are normally distributed, given all the model parameters, the joint distributions of  $(\mathbf{Y}_t^o, \mathbf{Y}_t^u)$  and  $(\mathbf{Z}_t^o, \mathbf{Z}_t^u)$  are multivariate normal. Then, the conditional distributions of  $\mathbf{Y}_t^u$  given  $\mathbf{Y}_t^o$  and  $\mathbf{Z}_t^u$  given  $\mathbf{Z}_t^o$  are also normal with moments of well-known form.

**5.3. Conditional for constant phase parameters**

The constant phase parameter  $\mathbf{a}$  only appears in the model equation for  $\mathbf{Y}_t$ . Given  $\boldsymbol{\theta}^y, \sigma_y^2$  and  $\lambda_y$ , it can be easily shown that the conditional likelihood for  $\mathbf{a}$  is given by the regression model

$$\mathbf{K}_t = a_1 \sin(\pi t/12) \boldsymbol{\alpha}_{1t} + a_2 \sin(\pi t/6) \boldsymbol{\alpha}_{2t} + \boldsymbol{\varepsilon}_t^y$$

where

$$\mathbf{K}_t = \mathbf{Y}_t - \mathbf{1}_n \beta_t^y - \cos(\pi t/12) \boldsymbol{\alpha}_{1t} - \cos(\pi t/6) \boldsymbol{\alpha}_{2t} - \mathbf{Z}_t \boldsymbol{\gamma}_t$$

is computed at the current value of  $\boldsymbol{\theta}^y$ . Then, the conditional likelihood function for  $\mathbf{a}$  has a bivariate normal kernel. Computations are then particularly easy if the prior is a conjugate bivariate normal or a standard reference prior  $p(\mathbf{a}) \propto 1$ .

**5.4. Spatial interpolation**

Let  $s$  denote an unobserved site;  $Y_t^s$  and  $Z_t^s$  are the unobserved square-root ozone concentration and temperature at such a site at time  $t$ . Spatial interpolation consists of producing samples of  $(Y_t^s, Z_t^s)$  from its posterior distribution. Obtaining such a sample is done by first sampling  $Z_t^s$  from

$$p(Z_t^s | \beta_t^z, \eta_t, \boldsymbol{\delta}_t, \sigma_z^2, \lambda_z, \mathbf{Z}_t),$$

and then sampling the parameters  $(\alpha_{1t}^s, \alpha_{2t}^s)$  that define the seasonal term at site  $\mathbf{s}$  from

$$p(\alpha_{it}^s | \alpha_{it-1}^s, \boldsymbol{\alpha}_{it}, \boldsymbol{\alpha}_{it-1}), \quad i = 1, 2.$$

Finally, we sample  $Y_t^s$  from

$$p(Y_t^s | \alpha_{1t}^s, \alpha_{2t}^s, Z_t^s, \beta_t^y, \boldsymbol{\alpha}_{1t}, \boldsymbol{\alpha}_{2t}, \gamma_t, \sigma_y^2, \lambda_y, \mathbf{Y}_t, \mathbf{Z}_t).$$

Since the state vector  $\boldsymbol{\theta}_z$  is independent of location, the conditional distribution of  $Z_t^s$  given  $\boldsymbol{\theta}_z$  and  $\mathbf{Z}_t$  is derived from the observation equation for  $(Z_t^s, \mathbf{Z}_t)$  and follows a normal distribution. To sample  $\alpha_{it}^s, i = 1, 2$ , we consider the augmented state vectors  $\boldsymbol{\alpha}_{it}^* = (\alpha_{it}^s, \boldsymbol{\alpha}_{it}), i = 1, 2$ , and derive its distribution from the corresponding evolution equation. Given a realization of  $(\alpha_{1t}^s, \alpha_{2t}^s)$ , the conditional for  $Y_t^s$  is derived from the joint distribution of the pair  $(Y_t^s, \mathbf{Y}_t)$ , as obtained from the observation equation. Note that the simulation that is needed for the spatial interpolation can be implemented after running the MCMC algorithm that uses only the data that are observed at the monitoring stations.

## 6. Results

We fitted our space–time model to the data that appear in Fig. 2 with the purpose of obtaining forecasts and kriged estimates of ozone concentrations. The hyperparameters of the prior distributions were chosen according to the fact that the evolution range is higher than the observation range. This is because harmonics are spatially more correlated than the errors in the observation equation, since the seasonality is quite similar across locations. Also, we assumed that the prior for the observation variance is larger for temperature than for ozone concentration, reflecting the fact that the temperature shows more variability than the square root of ozone concentration.

The results below are based on the following prior specifications. For the observation variogram parameters, we choose inverse gamma priors:  $\lambda_y \sim \text{IG}(1, 5)$  and  $\sigma_y^2 \sim \text{IG}(2, 0.01)$  for ozone concentration, and  $\lambda_z \sim \text{IG}(1, 0.5)$  and  $\sigma_z^2 \sim \text{IG}(1, 0.25)$  for temperature. For the evolution equation, we fix the range and variance parameters at the following values:  $\lambda_1 = 25, \lambda_2 = 25, \tau_1^2 = 0.02, \tau_1^2 = 0.0002, \tau_2^2 = 0.0004$  and  $\tau_\gamma^2 = 0.0002$  for ozone concentration, and  $\tau_\eta^2 = 18, \tau_z^2 = 0.004$  and  $\tau_\delta^2 = 0.04$  for temperature. The prior hyperparameters were selected on the basis of variogram fits to the univariate residuals. The evolution parameters were selected after an extensive simulation study, where we required the simulated data to match closely the spatial and seasonal patterns in the historical data.

We complete the specification with priors on the initial states. For ozone, we set  $\boldsymbol{\theta}_0^y \sim N(\mathbf{m}_0^y, \mathbf{C}_0^y)$  with  $\mathbf{m}_0^y = (2.85, -0.75\mathbf{1}_n', 0.08\mathbf{1}_n', 0.01)'$ , and  $\mathbf{C}_0^y = \text{block diag}(1, 0.01\mathbf{I}, 0.01\mathbf{I}, 0.001)$ , with the intercept centred on the mean square-root level, the harmonic parameters centred at values that match the bimodal diurnal cycle and the temperature coefficient centred at a slightly positive value. We choose a fairly large variance for all the components except for the harmonics. Finally, for temperature, we set  $\boldsymbol{\theta}_0^z \sim N(\mathbf{m}_0^z, \mathbf{C}_0^z)$ , with  $\mathbf{m}_0^z = (18.5, 0, -2.7, -5, 0, 0)'$ , and  $\mathbf{C}_0^z = \text{diag}(10, 1, 10, 10, 10, 10)$ , with the intercept centred at the September average, and the harmonics calibrated to capture the daily cycle.

Under these prior specifications, we ran the MCMC algorithm for 25000 iterations after a burn-in period of 1000 iterations, and we collected all samples for posterior inference. Table 1 reports the posterior median as well as the extremes for a 95% posterior probability interval for the parameters in the model which are not time varying.

**Table 1.** Posterior summaries for  $\sigma^2$ ,  $\lambda_y$ ,  $a_1$ ,  $a_2$ ,  $\sigma_z^2$  and  $\lambda_z$ 

Quantile	$\sigma_y^2$	$\lambda_y$	$a_1$	$a_2$	$\sigma_z^2$	$\lambda_z$
2.5%	1.15	5.60	2.14	7.43	0.81	5.76
Median	1.27	6.63	2.47	9.83	0.89	6.99
97.5%	1.40	7.84	2.85	13.26	0.98	8.37


Fig. 4 shows retrospective medians and predictive values of ozone level for three monitoring stations: Xalostoc, Benito Juárez and Pedregal. From September 8th until September 13th, medians and 95% probability bands (full curves) are plotted with the hourly ozone observations (dots). Also, we present the forecast median with the 95% predictive probability intervals (full curves) and the actual observed values of ozone level (open circles) for the 12 h of September 14th. In the retrospective sense, the model represents the cyclical patterns and non-stationarities of the data adequately. Also, the median and predictive intervals for September 14th look in accordance with the actual data. Note that our results are reported on the original scale and not in square-root units.

Maps of median ozone levels for September 8th, 1997, from 11 a.m. to 6 p.m. appear in Fig. 5. The maps are based on a  $15 \times 15$  interpolation grid over which we apply the spatial interpolation algorithm that is presented in Section 6. The median was computed across MCMC samples. The resulting map seems to be consistent with the cyclical behaviour of the data and theories about the dispersion of ozone in Mexico City. The pollutant builds at around 12 p.m., and the peak hours are between 2 and 4 p.m. The levels seem to dissipate by 6 p.m.

Fig. 6 shows the corresponding interquartile range maps for square-root ozone levels on September 8th. The maps illustrate three features of the model. First, the overall uncertainty level changes drastically from hour to hour from the peak uncertainty at 3 p.m. when the ozone levels are highest, to the lowest at 6 a.m. when the levels are lowest. Second, as expected, uncertainty is lowest near the observed stations. Finally, we note that the standard deviation contours are spherically shaped between stations, which is because of our assumption of isotropy for the errors.

Fig. 7 shows hourly means and intervals for temperature (in degrees Celsius) at the three stations over the 168-h period. For the first 144 h, we show the smoothed (retrospective) means and corresponding 95% intervals, based on the data up to time 144. For the last 24 h, we plot the predictive means and 95% intervals, again based on the first 144 h of data. The estimation of the altitude coefficient in the model for temperature reveals that there is a significant effect of such a variable, which is more pronounced during the night. From Fig. 7 and analogous plots for the rest of the stations, we observe that the model slightly underpredicts the maximum temperature of some stations, and overpredicts the minimum in others. We started our modelling by considering that the temperature was constant over all stations for each time step and found that the current model improves the fit of the ozone fields, which is our goal. We feel that further elaboration on the model for temperature will not produce results that are substantially better than those those which were obtained.

In Fig. 7(b), we show smoothed and predictive intervals at Benito Juárez, a station which does not take temperature measurements. This plot highlights a key feature of our modelling framework, namely that we can impute missing values and quantify uncertainty for either variable (temperature or ozone) at any specified location in the modelling domain.

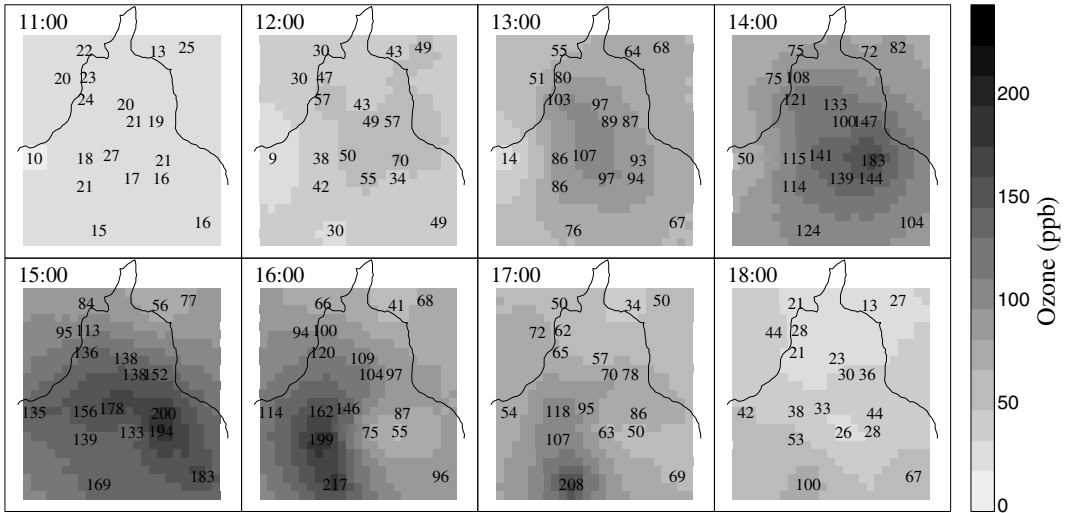


**Fig. 4.** Ozone data, retrospective medians and predictive values with 95% probability bands for three monitoring stations: (a) Xalostoc; (b) Benito Juárez; (c) Pedregal

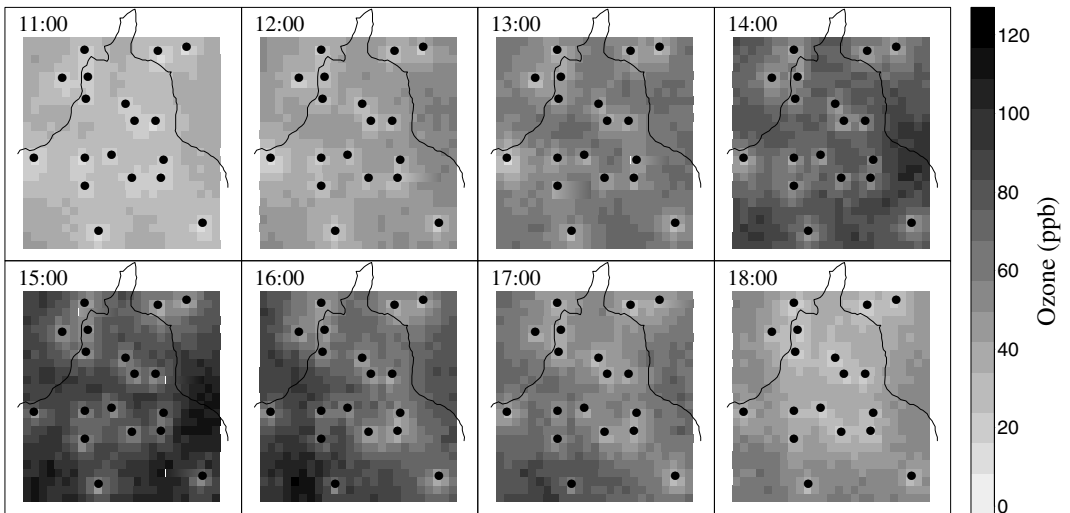
**6.1. Discussion**

A difficult task for spatiotemporal models with a non-trivial number of sites and time steps is to assess the validity of the model. We have not been able to find in the literature a comprehensive and systematic approach for spatiotemporal model validation. This is probably because of the difficulties in dealing with massive amounts of data and taking into account time and space. We concentrate our analysis on the study of the predictive behaviour of the model for which, given the Bayesian nature of our approach, we have a full probabilistic description.

To assess the spatial predictive capability of our model we took a ‘leave-one-out’ approach with each of the 19 stations. As an example, Fig. 8 shows the posterior predictive densities at 2 p.m. for September 10th, 1997. The labels on the *x*-axis correspond to the actual observed values, which were not used in the model fit. We observed that in most cases the actual observations fall in the central range of the corresponding predictive density, indicating compatibility between the model predictions and the observations. Additionally, the posterior forecast densities corresponding to 2 p.m. on September 14th, 1997, are shown in Fig. 9. In this case, the data from all the stations were used to compute the forecasts. Once again, the actual observed values




**Fig. 5.** Retrospective median ozone maps for September 8th, 1997, based on 25000 MCMC samples, using a 2.33 km<sup>2</sup> (15 × 15 pixel) mapping grid: the values that are written on the maps represent the observed data at that time and station



**Fig. 6.** Retrospective interquartile range maps of the square-root ozone concentrations for September 8th, 1997: ●, locations of the stations

fall within the range of the predictive densities, which confirms our comments in the previous section with respect to the forecast intervals that are presented in Fig. 4.

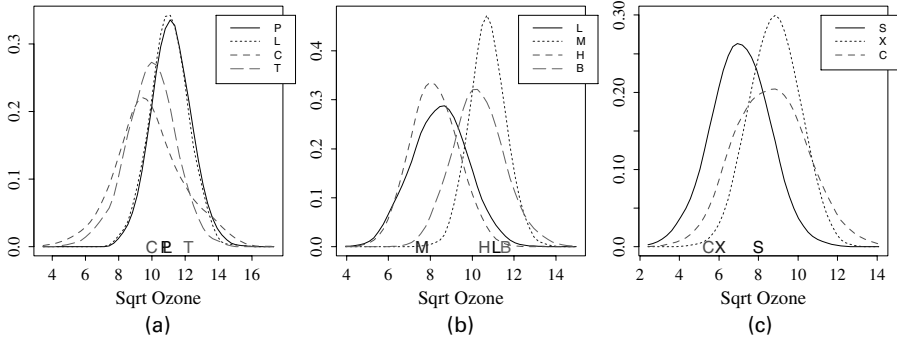
Kim *et al.* (1998) proposed a method to assess the validity of a time series model by using one-step-ahead predictions. We consider a modification of their method to the spatial case by considering the prediction that is produced by the model for a station that has been left out of the analysis. The idea is to compute the predictive probability that a ‘leave-one-out’ forecast is below its corresponding observed value. In other words, for an omitted station  $i$ , we obtained an estimate of  $u_{it} = \Pr(Y_{it}^* \leq Y_{it} | Y_{-i})$ , where the prediction is denoted by an asterisk and  $Y_{-i}$  corresponds to the data without station  $i$ , by averaging across MCMC iterations. It can be shown




**Fig. 7.** Temperature data, retrospective means and predictive values with 95% probability bands for three monitoring stations: (a) Xalostoc; (b) Benito Juárez; (c) Pedregal

that, if the model is correctly specified, the values in the sequence of  $u_{it}$  are independent and they are uniformly distributed on the interval  $(0, 1)$ . We explored the 19 possible sequences of  $u_{it}$ -values by using  $qq$ -plots and observed acceptable results in most cases. The worst behaviour was observed for Merced, for which the distributional assumptions were completely unsatisfied. A careful analysis of the data from this station reveals anomalously low values of ozone levels in relation to the neighbouring stations. We hypothesized a calibration problem, but, since similar discrepancies were not observed during the rest of the year, we suppose that there was either a transient failure in the measurement process or a very localized effect that our model cannot capture.

A by-product of the leave-one-out predictive analysis is the availability of draws from the posterior distributions of the parameters given by the 19 different data sets. These allowed us to perform a posterior robustness analysis with respect to the presence of each of the stations. The marginal posterior distribution of most parameters was fairly insensitive to the removal of a station. The exceptions were the range parameters  $\lambda_y$  and the phase parameter  $a_1$  for which box plots are shown in Fig. 10. We observed that removing Merced (MER) produced an important increase in  $\lambda_y$ . This is consistent with our previous comment regarding this station. In contrast,



**Fig. 8.** Posterior predictive densities at 2 p.m., September 10th, 1997, for the leave-one-out station analysis for three different geographical regions (the stations are identified with the first letter of their names; letters on the x-axis correspond to actual observations (which were not used in the predictions): (a) south-west region; (b) central region; (c) north-east region




**Fig. 9.** Posterior forecast densities at 2 p.m., September 10th, 1997, for three different geographical regions (the stations are identified with the first letter of their names; letters on the x-axis correspond to actual observations): (a) south-west region; (b) central region; (c) north-east region

removing Tlalpan (TPN) produced a significant decrease in  $a_1$ , something that is probably due to the downwind position of this station which sets it slightly out of phase with the others.

**6.2. Measurement error**

Given the anomalies that we observed in some stations, particularly Merced, we thought of including a specific component in the model to explain measurement error. Ozone measurement devices are usually calibrated to a precision of 3–5 ppb for, roughly, the whole measurable range. Our model considers the data after a square-root transformation; we thus notice that a first-order Taylor expansion yields  $\sqrt{(\mu + v)} \approx \sqrt{\mu} + 1/(2\sqrt{\mu})v$  where  $\mu$  is the mean level and  $v$  the measurement error. We assume that a measurement of 5 ppb corresponds to 1 standard deviation of each of the minute-by-minute readings of the instrument. Since our observations correspond to 1-h averages and the smallest reading amounts to 3 ppb, an upper bound of the standard deviation of the error term in the previous equation is  $5/2\sqrt{(3 \times 3600)} = 0.02406$ . So, under the assumption that the measurement errors are uncorrelated and that there is no bias due to miscalibration of the instruments, the measurement error is negligible.

A related issue is that of incorporating a ‘nugget’ effect in the model to capture small range variabilities. This can be done by adding an extra random term or, equivalently, a third layer to



**Fig. 10.** Box plots of samples from the marginal posterior distributions of (a)  $\lambda_\gamma$  and (b)  $a_1$  obtained by leaving one station out (indicated on the horizontal axis); the plot denoted ALL represents the posterior sample when all stations are included

the current model. We feel that our model already incorporates random effects for location, temperature and periodicity. A nugget effect will produce predictions that are less smooth, which is probably realistic but at the cost of dealing with identifiability problems and adding parameters to an already complex model. Cressie and Wikle (1998) advocated the use of a three-layer model which focuses on a direct specification of the physical dynamics. This is our preferred path for elaborations over the present approach.

### 7. Conclusions and extensions

In this paper, we introduced a new spatiotemporal model for hourly ozone levels in Mexico City that allows for spatial interpolation and prediction. The model is formulated within the state space framework and includes a set of time-varying Fourier coefficients to account for the periodicity in the data. Also, the model considers uncertainty on any missing values of ozone concentrations and the covariates which, in the present version, is only the temperature. We produced posterior inference with an efficient block MCMC simulation algorithm that samples the spatial covariance and error variance parameters unconditional on the state vectors. The



results that were obtained with the model seem reasonable for short-term forecasting and spatial interpolation, as can be assessed from the predictive analysis.

However, the actual model is only empirical and does not consider transport or chemical reactions that are related to ozone. Currently, we are trying to obtain second-by-second data, wind information and other covariables such as NO<sub>x</sub> or volatile organic compounds to incorporate in the model. We believe that this extra information combined with scientific priors based on climatology or photochemistry can produce a space–time model that not only predicts and interpolates adequately but also has some physical interpretation.

**Acknowledgements**

We thank Cristina Ortuño of Departamento del Distrito Federal in Mexico City for providing the data that were analysed in this paper. We also thank Michael Stein, John Frederick and Jeff Gaffney for their helpful comments and suggestions. G. Huerta and B. Sansó were partially funded by the Consejo Nacional de Ciencia y Tecnología (Mexico)–Consejo Nacional de Investigaciones Científicas y Tecnológicas (Venezuela) international exchange grant PI-2000000834. J. Stroud was funded by National Science Foundation group infrastructure grant 97-09696 while he was at the University of Chicago. This work was also partially funded by Consejo Nacional de Investigaciones Científicas y Tecnológicas grant G97-000592 and the Consejo Nacional de Ciencia y Tecnología grant J34413-E.

**Appendix A: Marginal likelihood for range and variance parameters**

The models that are presented in this paper have the general DLM or state space form where

$$p(\mathbf{Y}_t|\boldsymbol{\theta}_t, \sigma^2, \lambda) = N(\mathbf{Y}_t|\mathbf{F}'_t\boldsymbol{\theta}_t, \sigma^2\mathbf{V}_\lambda),$$

$$p(\boldsymbol{\theta}_t|\boldsymbol{\theta}_{t-1}) = N(\boldsymbol{\theta}_t|G\boldsymbol{\theta}_{t-1}, \sigma^2\mathbf{W}_\lambda),$$

with  $\boldsymbol{\theta}_t$  denoting the state vector and  $\lambda$  a non-negative parameter that defines the covariances  $\mathbf{V}_\lambda$  and  $\mathbf{W}_\lambda$ , and  $\sigma^2$  a common scale factor,  $p(\boldsymbol{\theta}_0|D_0) = N(\boldsymbol{\theta}_0|\mathbf{m}_0, \mathbf{C}_0)$ . Denoting  $D_t$  the information that is available until time  $t$  and assuming that

$$p(\lambda, \sigma^2) = p(\lambda) p(\sigma^2)$$

then the joint posterior distribution for the state vectors, the variance and the range parameter is

$$p(\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_T, \sigma^2, \lambda|D_T) = p(\lambda) p(\sigma^2) p(\boldsymbol{\theta}_0|D_0) \prod_{t=1}^T p(\mathbf{Y}_t|\boldsymbol{\theta}_t, \sigma^2, \lambda) p(\boldsymbol{\theta}_t|\boldsymbol{\theta}_{t-1}, \sigma^2, \lambda).$$

Using Bayes’s theorem and the Markovian structure of the model, the joint posterior can be written as

$$p(\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_T, \sigma^2, \lambda|D_T) = p(\lambda) p(\sigma^2) \prod_{k=1}^T p(\boldsymbol{\theta}_{T-k}|\boldsymbol{\theta}_{T-k+1}, \sigma^2, \lambda, D_T) \prod_{t=1}^T p(\mathbf{Y}_t|\sigma^2, \lambda, D_{t-1})$$

$$= p(\boldsymbol{\theta}|\sigma^2, \lambda, D_T) p(\sigma^2|\lambda, D_T) p(\lambda|D_T)$$

where  $\boldsymbol{\theta} = (\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_T)$ .

Given that  $p(\mathbf{Y}_t|\sigma^2, \lambda, D_{t-1}) = N(\mathbf{f}_t, \sigma^2\mathbf{Q}_t)$ , where  $\mathbf{f}_t$  and  $\mathbf{Q}_t$  are obtained with the forward filtering equations in West and Harrison (1997), chapter 16. Integrating out  $\boldsymbol{\theta}$  we obtain that the posterior of  $(\lambda, \sigma^2)$  can be expressed as

$$p(\lambda, \sigma^2|D_T) \propto p(\lambda) p(\sigma^2) \prod_{t=1}^T p(\mathbf{Y}_t|\sigma^2, \lambda, D_{t-1})$$

$$\propto p(\lambda) p(\sigma^2) \prod_{t=1}^T |\mathbf{Q}_t|^{-1/2} \left(\frac{1}{\sigma^2}\right)^{T/2} \exp\left\{-\frac{1}{2\sigma^2} \sum_{t=1}^T (\mathbf{Y}_t - \mathbf{f}_t)' \mathbf{Q}_t^{-1} (\mathbf{Y}_t - \mathbf{f}_t)\right\}.$$

Thus, if we adopt an inverse gamma prior with parameters  $a_\sigma$  and  $b_\sigma$  for  $\sigma^2$ , we have that  $p(\sigma^2|\lambda, D_T)$  is an inverse gamma distribution with shape parameter  $T/2 + a_\sigma$  and scale parameter  $\sum_{t=1}^T (\mathbf{Y}_t - \mathbf{f}_t)' \mathbf{Q}_t^{-1} (\mathbf{Y}_t - \mathbf{f}_t) / 2 + b_\sigma$ .

Additionally, if we integrate out  $\sigma^2$  with this prior, then

$$p(\lambda|D_T) \propto p(\lambda) \prod_{t=1}^T |\mathbf{Q}_t|^{-1/2} \left\{ \sum_{t=1}^T (\mathbf{Y}_t - \mathbf{f}_t)' \mathbf{Q}_t^{-1} (\mathbf{Y}_t - \mathbf{f}_t) + b_\sigma \right\}^{-T/2 + a_\sigma}.$$

Both  $\mathbf{Q}_t$  and  $\mathbf{f}_t$  depend on  $\lambda$  since the forward filtering equations involve  $\mathbf{V}_\lambda$  and  $\mathbf{W}_\lambda$ .

Finally, since

$$p(\boldsymbol{\theta}|\sigma^2, \lambda, D_T) = \prod_{k=1}^T p(\boldsymbol{\theta}_{T-k}|\boldsymbol{\theta}_{T-k+1}, \sigma^2, \lambda),$$

and each of the terms in the product is a normal distribution, a sample of  $\boldsymbol{\theta}$  can be obtained recursively moving backwards from  $\boldsymbol{\theta}_T$  to  $\boldsymbol{\theta}_1$ .

## References

- Berliner, M., Royle, A., Wikle, C. and Milliff, R. (1999) Bayesian methods in the atmospheric sciences. In *Bayesian Statistics 6* (eds J. M. Bernardo, J. O. Berger, A. P. Dawid and A. F. M. Smith), pp. 83–100. Oxford: Oxford University Press.
- Bretthorst, L. (1988) *Bayesian Spectral Analysis and Estimation*. New York: Springer.
- Carroll, R., Chen, R., George, E., Li, T., Newton, H., Schmiediche, H. and Wang, N. (1997) Ozone exposure and population density in Harris County, Texas. *J. Am. Statist. Ass.*, **92**, 392–415.
- Carter, C. and Kohn, R. (1994) Gibbs sampling for state space models. *Biometrika*, **81**, 541–553.
- Cressie, N. and Wikle, C. (1998) Strategies for dynamic space-time statistical modeling: discussion of The kriged kalman filter by Mardia, Goodall, Redfern and Alonso. *Test*, **7**, 257–264.
- Frühwirth-Schnatter, S. (1994) Data augmentation and dynamic linear models. *J. Time Ser. Anal.*, **15**, 183–202.
- Guttorp, P., Meiring, W. and Sampson, P. (1994) A space-time analysis of ground-level ozone data. *Environmetrics*, **5**, 241–254.
- Kim, S., Shephard, N. and Chib, S. (1998) Stochastic volatility: likelihood inference and comparison with arch models. *Rev. Econ. Stud.*, **65**, 361–393.
- Knorr-Held, L. and Rue, H. (2002) On block updating in Markov random field models for disease mapping. *Scand. J. Statist.*, **29**, 597–614.
- Mardia, D., Goodall, C., Redfern, E. and Alonso, F. (1998) The kriged kalman filter (with discussion). *Test*, **7**, 217–285.
- Meiring, W., Guttorp, P. and Sampson, P. (1998) Space-time estimation of grid-cell hourly ozone levels for assessment of a deterministic model. *J. Environ. Ecol. Statist.*, **5**, 197–222.
- Milanchus, M., Rao, T. and Zurbenko, I. (1998) Evaluating the effectiveness of ozone management efforts in the presence of meteorological variability. *J. Air Waste Mangmnt Ass.*, **48**, 201–215.
- Rao, S., Zurbenko, I., Neagu, R., Porter, P., Ku, J. and Henry, R. (1997) Space and time scales in ambient ozone data. *Bull. Am. Meteorol. Soc.*, **78**, 2153–2166.
- Rue, H. (2001) Fast sampling of Gaussian Markov random fields. *J. R. Statist. Soc. B*, **63**, 325–338.
- Sampson, P. and Guttorp, P. (1992) Nonparametric estimation of nonstationary spatial covariance structure. *J. Am. Statist. Ass.*, **87**, 108–119.
- Sansó, B. and Guenni, L. (2000) A non-stationary multisite model for rainfall. *J. Am. Statist. Ass.*, **95**, 1064–1089.
- Shaddick, G. and Wakefield, J. (2002) Modelling daily multivariate pollutant data at multiple sites. *Appl. Statist.*, **51**, 351–372.
- Stroud, J. R., Müller, P. and Sansó, B. (2001) Dynamic models for spatiotemporal data. *J. R. Statist. Soc. B*, **63**, 673–689.
- Thompson, M., Reynolds, J., Cox, L. and Guttorp, P. (1999) A review of statistical methods for the meteorological adjustment of tropospheric ozone. *Technical Report*. University of Washington, Seattle.
- Tonellato, S. (1997) Bayesian dynamic linear models for spatial time series. *Rapporto di Ricerca 5/1997*. Dipartimento di Statistica, Università Ca' Foscari di Venezia, Venice.
- West, M. and Harrison, P. (1997) *Bayesian Forecasting and Dynamic Models*, 2nd edn. New York: Springer.
- Wikle, C., Berliner, M. and Cressie, N. (1999) Hierarchical Bayesian space-time models. *J. Environ. Ecol. Statist.*, **5**, 117–154.