

Conditionally specified space-time models for multivariate processes

Michael J. Daniels
Department of Statistics
University of Florida
Gainesville, FL 32606 USA
mdaniels@stat.ufl.edu

Zhigang Zhou
Department of Statistics
Iowa State University
Ames, IA 50011 USA
zgzhou@iastate.edu

Hui Zou
Department of Statistics
Stanford University
Stanford, CA 94305-4065 USA
hzou@stat.stanford.edu

Summary

We propose a class of conditionally specified models for the analysis of multivariate space-time processes. Such models are useful in situations where there is sparse spatial coverage of one of the processes and much more dense coverage of the other process(es). The dependence structure across processes and over space, and time is completely specified through a neighborhood structure. These models are applicable to *both* point and block sources; for example, multiple pollutant monitors (point sources) or several county level exposures (block sources). We introduce several computational tricks which are integral for model fitting, give some simple sufficient and necessary conditions for the space-time covariance matrix to be positive definite, and implement a Gibbs sampler, using Hybrid MC steps, to sample from the posterior distribution of the parameters. Model fit is assessed via the DIC. Predictive accuracy, both over time and space, is assessed both relatively and absolutely via mean squared prediction error and coverage probabilities. As an illustration of these models, we fit them to particulate matter and ozone data collected in the LA area in 1995 over a 3 month period. In this data, the spatial coverage of particulate matter was sparse relative to that of ozone.

1 Introduction

Conditionally specified, or Markov random field, models can often be easily specified for processes over space and time, but computations are intractable. We propose a particular class of conditionally specified space-time models for a multivariate process here that will be computationally tractable and that allows a neighborhood structure to be explicitly specified over both space and time. These models will also accommodate nonstationary spatial and temporal structures and are applicable to both point and block source data, whether as a model for the data itself or as a model for random effects embedded within a regression model for the data.

Some recent work on conditional specification of models for spatial data include Cressie and Chan (1989) and Cressie et al. (1999) who use 'distance'-based neighbors for block and point sources, respectively. Bayesian conditional autoregressive (CAR) models using distance-based neighbors have been applied in Conlon and Waller (1998) and Best et al. (1999). Kaiser et al. (2002) specify anisotropic distance-based neighborhood structures which vary over time. Conditions under which conditionally specified models induce a valid joint distribution were originally given in Besag (1974) and a weaker set of conditions was given in Kaiser and Cressie (2000). Hrafnkelsson and Cressie (2003) compare geostatistics and conditionally specified models for predicting at many location using a fine grid; they conclude that the conditional approach is a viable alternative for spatial environmental data. Recent Bayesian work on specifying priors for particular conditionally specified models includes a series of papers by Sun, Tsutakawa, and colleagues (Sun, Tsutakawa, and Speckman, 1999; Kim, Sun, and Tsutakawa, 2001); these papers and many papers using CAR models focus on placing these conditional priors on random effects in the setting of block (areal) sources.

There has also been work on conditionally specified models for **multivariate** spatial processes. Mardia (1988) generalizes CAR models to the multivariate setting, referring to them as MCAR models. These models are somewhat restrictive as the same spatial

dependence parameters are assumed for each of the multivariate processes. Gelfand and Vounatsou (2003) generalize this work to allow a separate spatial dependence parameter for each multivariate process. Carlin and Banerjee (2003) propose separable models with a CAR model for spatial dependence and direct specification of the multivariate dependence. Knorr-Held and Rue (2002) develop intrinsic CAR models for bivariate processes.

In addition, there has also been considerable work on space-time modeling in recent years, in some cases with multivariate processes. Carroll et al. (1997) specify a space-time covariance function for a Gaussian random field on ozone (for a discussion of their choice, see the accompanying discussion of this paper by Cressie). Wikle, Berliner, and Cressie (1998) build complex Bayesian space-time models that capture spatial dependence using Gaussian Markov random fields (GMRFs) and temporal dependence using vector autoregressive (VAR) models. Wikle and Cressie (1999) develop a space-time Kalman filter that is more general than the (kriged) Kalman filter previously proposed in Mardia et al. (1998). Le et al. (1997) and Kibria et al. (2002) use multivariate t-models for daily multivariate spatial prediction of pollutants while assuming independence over time using the concept of gauged (sometimes observed) and ungauged (never observed) sites. To deal with the space-time problem, Zidek et al. (2002) extends this approach by taking temporal correlation into account through pre-whitening the time series (as a block) which results in minimal loss of the spatial correlation of the residuals, which are again modeled using multivariate t-models. Haas (1995, 2000) uses a moving “cylinder” to do local prediction in a space time process. Niu et al. (2003) propose a general model for spatio-temporal data on a regular lattice; after ‘differencing’ to remove seasonal patterns and/or trends, they regress the response at each location on the grid at each time using ARMA models to account for temporal dependence and order r spatial neighborhoods. Stein (2003) recently proposed a class of non-separable space-time covariance functions which require pseudo-likelihood approaches for estimation. Shaddick and Wakefield (2002) construct models for multivariate space-time models using geostatistics models for the spatial component and vector autoregressive models for the temporal and

multivariate dependence. Other work in the development of space-time models includes Handcock and Wallis (1994), Stroud et al. (2001), and Brown et al. (2001).

In general, computations in space-time models are often problematic, requiring the inversion of large covariance matrices and as such, separable (over space and time) models are often considered (Gelfand et al., 2001). In addition, when both spatial and temporal dependence are considered the literature is lacking on models with such dependence represented explicitly via spatial *and* temporal neighbors. Some work in this area can be found in Knorr-Held (2000) who considers neighborhoods among spatial and temporal random effects and with space-time interaction terms.

A general issue these models will address is the use of information on a 'densely' observed space-time processes to help 'fill-in' a correlated, but sparsely observed, space-time process, whether the process is observed at a point or block. As an illustration of this issue and the methodology, we analyze data from a joint pollutant field of particulate matter and ozone in the Los Angeles (LA) area (van den Eeden, 2001), where there were many monitors for ozone, but only a few for particulate matter. Models such as those proposed by Le et al. (1997) and Zidek et al. (2002) would be well suited to this application, but would not generalize to block source data for which conditionally specified models such as the one proposed here are better suited.

Section 2 will introduce the model and discuss specification of the components of the covariance structure, computations, and model fitting. Section 3 presents the results from analyzing the LA pollutant data. Some extensions of the model are proposed in Section 4. Conclusions and discussion are contained in Section 5.

2 Model

Let $Y_k(s_{ki}, t) : k = 1, \dots, K; i = 1, \dots, N_k; t = 1, \dots, T$ denote the possibly transformed measured value for the k th process (pollutant) at its i th location at time t , with $p = \sum_k N_k$ the total number of sub-sites. We refer to the s_{ki} as sub-sites since more than one of the K

processes may be measured at the same location. Now, consider models with the following form:

$$Y \sim N(W\beta, [M^{-1}(\tau^2)(I_{pT} - C(\eta))]^{-1}) \quad (1)$$

where $Y = (Y^*(1)', Y^*(2)', \dots, Y^*(T)')$ is a $pT \times 1$ vector, with

$$Y^*(t) = (Y_1(s_{11}, t), Y_1(s_{12}, t), \dots, Y_1(s_{1,N_1}, t), Y_2(s_{21}, t), \dots, Y_2(s_{2,N_2}, t), \dots, Y_K(s_{K,N_K}, t))$$

a $p \times 1$ vector, M is a block diagonal matrix of the form, $M = I_T \otimes M^*$, where M^* is a $p \times p$ matrix, and

$$C(\eta) = \begin{pmatrix} V_s(\eta_s) & V_t(\eta_t) & 0 & \dots \\ V_t(\eta_t) & V_s(\eta_s) & V_t(\eta_t) & \dots \\ 0 & V_t(\eta_t) & V_s(\eta_s) & \dots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix}. \quad (2)$$

$V_s(\eta_s)$ is a $p \times p$ matrix characterizing the spatial neighborhoods with $\text{diag}(V_s) = 0_p$ and parameters η_s , and $V_t(\eta_t)$ is a $p \times p$ matrix, characterizing the temporal neighborhoods, parameterized by η_t .

This covariance structure implies the following two conditional specifications, the first in terms of an individual element of the Y vector when M^* is assumed to be diagonal, and the second in its partitioned form, $Y^*(t)$, where M^* is not assumed diagonal,

$$Y_k(s_{ki}, t) | Y_{k'}(s_{k'i'}, t') \neq Y_k(s_{ki}, t) \sim N(w_{ikt}\beta + \sum_{i',k',t'} c_{(i,i'),(t,t')}^{(k,k')} (y_{k'}(s_{k'i'}, t') - w_{i'k't'}\beta), \tau_{ik}^2) \quad (3)$$

$$Y^*(t) | Y^*(j), j \neq t \sim N(W_{..t}\beta + V_t(\eta_t)[(Y_{t-1}^* - W_{..t-1}) + (Y_{t+1}^* - W_{..t+1})], M^*) \quad (4)$$

where $c_{(i,i'),(t,t')}^{(k,k')}$ denotes the element of the $C(\eta)$ matrix corresponding to the pair of values $(Y_k(s_{ki}, t), Y_{k'}(s_{k'i'}, t'))$, w_{ikt} is the row of the design matrix W corresponding to the i th location for the k th process at time t , $W_{..t}$ are the rows of the design matrix corresponding to time t , and τ_{ik}^2 is the diagonal element from the M matrix corresponding to $Y_k(s_{ki}, t)$. Note that these models are often constructed starting with conditional specification (3) or (4).

Some features and remarks on the model:

- These models assume a constant spatial neighborhood and parameters over time. However, there are no restrictions on the spatial neighborhood matrix, V_s (other than $M^*^{-1}(I_p - V_s)$ being positive definite).
- These models assume an AR(1)-type temporal neighborhood structure through the matrix $V_t(\eta_t)$. However, nonstationarity is allowed by varying the AR(1) parameters over space, and different sites can be lag 1 neighbors. The elements of this matrix are related to the coefficients of vector autoregressive (VAR) models (they are, in fact, just appropriately re-scaled versions of these parameters).
- Proper spatial CAR (e.g., Best et al, 1999) and MCAR (Mardia, 1988) models are special cases when $V_t = 0_p$ and the appropriate structure is placed on M^* and V_s . The η_s parameters play the role of the spatial autocorrelation parameters in proper spatial (M)CAR models.
- The conditional specifications of this model (3) or (4) demonstrate how straightforward predictions using both spatial and temporal neighbors can be done.
- Although the models have a similar form to separable geostatistics models (see, e.g., Gelfand et al., 2001), the resulting covariance matrix in these models is **not** separable.
- Asymptotics in these models when modeling point sources are not well understood. In Kaiser et al. (2002), no measures of uncertainty are reported for the parameter estimates. We will fit fully Bayesian models here to conduct exact Bayesian inference. This will facilitate taking into account all sources of variability when characterizing the variability of the predictions and assessing the variability of the parameters, in general.

2.1 Neighborhood structure, dependence parameters, and conditional (co-)variances

The neighborhood/dependence structure given in the spatial matrix, V_s , and the temporal matrix, V_t , can be specified with great flexibility subject to the constraint that $M^{-1}(I_{pT} - C(\eta))$ is positive definite. For example, if a natural barrier exists between two sites, the corresponding component of $C(\eta)$ can be set to 0. Such issues can be difficult to handle with direct specification of covariances, as in geostatistics models. Choices for the M^* matrix and the spatial and temporal neighborhood matrices, V_s and V_t , respectively, can be made by equating these matrices to their counterparts in common (multivariate) spatial models and temporal models. For example, if $N_k = N$ for $k = 1, \dots, K$ and $s_{ki} = s_{k'i}$ for all i , then the form of the M^* and V_s can be specified as in MCAR models (Mardia, 1988). Extending this to the more flexible MCAR models in Gelfand and Vounatsou (2003) will be discussed in Section 4.3. Any type of spatial dependence/parameters can be used in $V_s(\eta)$ as long as they are constant over time. The form of the V_t matrix can be made similarly to the choices commonly made in vector autoregressive (VAR) time series models for multivariate responses. It might be chosen to be diagonal with a separate parameter for each process/site or to be a full matrix with the temporal autoregressive parameters decreasing with distance from the site. Simplified (diagonal) forms for the M^* matrix include $M^* = \tau^2 I_p$, a constant conditional variance across sites and time and process or $M^* = \text{diag}_p(\tau_k^2)$, i.e., a separate conditional variance for each process. Additional discussion of specification of the M^* matrix can be found in Section 4.4. We will introduce the neighborhood structure, η parameters, and the parameterization of the M^* matrix used in our example in Section 3.

2.2 Prior distributions

A diffuse normal prior was assumed on the regression parameters, β , and inverse Wishart/gamma priors on M^* matrix and/or its components as appropriate. Prior specification for the dependence parameters, η , is tricky and an open problem. Here, for simplicity, we consider

uniform priors on the space over which $M^{-1}(I_{pT} - C(\eta))$ is positive definite. For some models, this region can be determined explicitly. For details, see Section 3.3. Prior specification on simpler CAR models within a multilevel setup can be found in Sun et al. (1999).

2.3 Computational Difficulties and Solutions

Space-time models often create computational challenges. Here, we will discuss how this class of models can be fit with minimal computational difficulty. The main challenge involves computation of both the determinant, inverse, and eigenvalues of a matrix which is $Tp \times Tp$ where T is the number of time points and $p = \sum_k N_k$ is the total number of spatial locations over each process. To accomplish these tasks, clever, but simple, computational approaches are required. We consider the case where $M = I_T \otimes M^*$.

Result 1:

If $M^{*-1}(I_p - V_s)$ is not positive definite, then $M^{-1}(I_{pT} - C(\eta))$ is not positive definite.

Proof:: It is a leading submatrix of $M^{-1}(I_{pT} - C(\eta))$.

Result 2:

Given the form of $M^{-1}(I_{pT} - C(\eta))$, we can rewrite it as

$$M^{-1}(I_{pT} - C(\eta)) = [I_T \otimes M^{*-1}(I_p - V_s(\eta_s))] - \rho \otimes (M^{*-1}V_t(\eta_t)) \quad (5)$$

where ρ is the $T \times T$ matrix defined below.

$$\rho = \begin{pmatrix} 0 & 1 & 0 & \dots \\ 1 & 0 & 1 & \dots \\ 0 & 1 & 0 & \dots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix}.$$

Result 3: Using some Kronecker product identities,

$$M^{-1}(I_{pT} - C(\eta)) = (I_T \otimes A^{-1})(I_{pT} - \rho \otimes B)(I_T \otimes A^{-1})$$

where $A = (I_p - V_s)^{-1/2}M^{*1/2}$ and $B = AM^{*-1}V_tA$. This result will hold given Result 1.

Result 4: $M^{-1}(I_{pT} - C(\eta))$ is positive definite if and only if A is full rank and $(I_{pT} - \rho \otimes B)$ is positive definite. (follows from an adaptation of Theorem A.1.1 in Anderson (1984, p. 583)).

Theorem I:

Define γ_i to be the eigenvalues of $M^{*-1}(I_p - V_s)$, λ_i to be the eigenvalues of ρ , and μ_j to be the eigenvalues of B . Then $M^{-1}(I_{pT} - C(\eta))$ will be positive definite iff $\gamma_i > 0$ and $\lambda_i \mu_j < 1$ for all i and j .

Proof: Follows directly from Results 1-4.

So to check the positive definiteness of the joint covariance matrix $M^{-1}(I_{pT} - C(\eta))$, we only need to compute the eigenvalues of two $\sum_k N_k$ -dimensional matrices, A and B , and a T -dimensional matrix ρ , which has a simple form and is fixed; this implies that even within a sampling based algorithm, the eigenvalues of ρ only need to be computed initially. Given the eigenvalues, it is easy to compute the determinant. To invert $M^{-1}(I_{pT} - C(\eta))$, we just need the eigenvectors of these same matrices. Thus, all the computations reduce to dealing with T and $p = \sum_k N_k$ dimensional matrices. Based on specific neighborhood structure, calculations can be simplified even further (see Section 3.3).

The form of the $C(\eta)$ matrix given in equation (2) implies a 'balanced' $C(\eta)$ matrix, i.e., the same sites over time and the same times for all sites. The computational tricks above require such balance. If the actual data is unbalanced, balance can be achieved by filling in the intermittent 'missing' data using data augmentation within an MCMC algorithm. The details of this will be discussed in Section 2.5.

So, similar to separable space-time geostatistics models, the 'limiting' factor in computations is the size of T and $p = \sum_{k=1}^K N_k$. Manipulating other T , p , or N_k dimensional matrices can sometimes facilitate posterior computation (see Section 3.3).

2.4 Posterior sampling

We use a Gibbs sampler to obtain a sample from the posterior distribution of the parameters. The full conditional distributions for β is normal. Sampling M^* will depend on the form of this matrix; e.g., if $M^* = \tau^2 I_p$, then the full conditional distribution of τ^2 will be inverse gamma. Otherwise, we propose to use an appropriate version of the Metropolis-Hastings algorithm. To sample the η , we used a Hybrid MC algorithm (Neal, 1994). This requires evaluation of the first derivative of the log full conditional with respect to η , which is simple to compute given the form of the likelihood,

$$\log(L(\eta)) \propto (1/2) \log |M^{-1}(I_{pT} - C(\eta))| - (1/2) \text{tr}[M^{-1}(I_{pT} - C(\eta))(Y - W\beta)(Y - W\beta)^T] \quad (6)$$

as η appear directly in the components of the $C(\eta)$ matrix. This algorithm facilitates efficient movement through the multivariate η space.

When intervals over which the η resulted in a positive definite $M^{-1}(I_{pT} - C(\eta))$ matrix could be determined explicitly, the parameters were transformed using a logit type transformation on that interval, $\log(\frac{\eta - \eta_{min}}{\eta_{max} - \eta})$, where (η_{min}, η_{max}) is the interval; this was done in order to transform η from a closed interval for the Hybrid MC algorithm. If an interval cannot be determined, a check of positive definiteness (p.d) can be conducted using Theorem I, and if the values corresponds to a non-positive definite covariance matrix, the value is thrown out (since it has prior probability 0). Details using the structure of the C and M matrix in our example will be given in Section 3.3.

2.5 Model comparison

As discussed in Kaiser et al. (2002), evaluation of model fit and predictive ability are two different questions and can result in differing conclusions. To evaluate model fit and compare models, we use the Deviance information criterion (DIC) (Spiegelhalter et al., 2002). We will give some details in the following. Define the parameter vector to be θ . The DIC is

defined to be $Dev(\hat{\theta}) + 2p_D$, where $Dev(\theta) = -2\loglik(\theta)$, $\hat{\theta}$ is the posterior mean of θ , and $p_D = \hat{Dev} - Dev(\hat{\theta})$ with \hat{Dev} the posterior mean of the Deviance. Here, $\theta = (\beta, M(\tau^2), \eta)$ and the likelihood is given in (6). Smaller values of the DIC are preferred.

To compare the predictive ability of these models, we compute both the prediction error and coverage probabilities. We use a cross-validation approach. To assess spatial prediction, we drop out the data on a sub-site at all times, re-fit the model, and then predict the process values at the dropped sub-site. For temporal prediction, we drop out all sub-sites at the first and last times, re-fit the model, and then predict the process values at the dropped times.

Some details on prediction follow (under the restriction that M^* is diagonal; extension to non-diagonal M^* is straightforward). Similar to the approach in Kaiser et al. (2002), we drop a sub-site (time) and fit the model to the remaining sites to obtain the posterior distributions of the parameters conditional on these sub-sites. Then, we assume the pollution process at the dropped sub-site follows the same distributional form as (3), i.e., the conditional distribution of Y_0 , the vector of responses at the sub-sites/times to be predicted (i.e., the dropped sub-sites/times), given the remaining sub-sites/times, Y_{obs} , follows a normal distribution with mean,

$$Y_{0,pred} = W_0\beta + \sum_{k=1}^K \sum_{i=1}^{N_k} \sum_{t=1}^T c_{(0,i),(0,t)}^{(0,k)} (Y_k(s_{ki}, t) - W_{ikt}\beta) \quad (7)$$

and variance

$$V_y^{-1} = M_0^{-1}(I_{dim(Y_{0,pred})} - C_0) \quad (8)$$

where C_0 is the neighborhood matrix for the sub-sites to be predicted, W_0 is the design matrix for the sub-sites to be predicted, $c_{(0,i),(0,t)}^{(0,k)}$ is the matrix of elements of the $C(\eta)$ matrix corresponding to the regression of the sub-sites to be predicted on the remaining sub-sites (cf. (3)), and M_0 is a diagonal matrix containing the appropriate conditional variances; these are formed using (3). So, for prediction, the dimension of the matrix that needs to be inverted will be the dimension of Y_0 .

The $C(\eta)$ matrix is constructed to include intermittent missing values. At each iteration of the sampler, these values are filled in via data augmentation using a normal distribution with mean and variance analogues of (7) and (8); this is necessary for the computational results outlined in Section 2.3 to hold. We emphasize that by using this approach, we are implicitly using monte carlo integration to account for the uncertainty of these missing values.

We will evaluate prediction using mean squared prediction error (MSPE) and coverage probabilities. Mean square prediction error is defined as

$$MSPE = \sum (Y_{0,pred} - Y_{0,obs})^2$$

where $Y_{0,pred}$ is the conditional posterior mean given in (7) and $Y_{0,obs}$ is the observed value of the dropped sub-site/time. To help calibrate MSPE, we can compare its value to the conditional variances (τ_{ik}^2), estimated from fitting the model with all sub-sites/times.

To evaluate coverage, we compute 95% prediction intervals using the following approach. We will illustrate the approach under the setting of dropping out a sub-site at all times (dropping out all sub-sites at one time follows similarly). At iteration m (of M) of the Gibbs sampler, we create 95% credible regions using $(Y_{0,obs}(s_i, \cdot) - Y_{0,pred}^m(s_i, \cdot))^T (V_y^m)^{-1} (Y_{0,obs}(s_i, \cdot) - Y_{0,pred}^m(s_i, \cdot))$ where $Y_{0,pred}^m(s_i, \cdot)$ is given in (7) and V_y^m is given in (8), using the values of (τ^2, η, β) at iteration m . Since the conditional distribution is normal, this should follow a chi-squared distribution on T degrees of freedom. We can count for how many iterations these regions cover the true value of the dropped sub-site; ideally, close to 95% for each sub-site. A referee has pointed out that these regions may be too small, especially if there is considerable uncertainty/variability reflected in the posterior distribution of (β, η, τ^2) . As an alternative, we also examine marginalized credible regions that integrate over the uncertainty in (β, η, τ^2) . These are constructed as $(Y_{0,obs}(s_i, \cdot) - Y_{0,pred}(s_i, \cdot))^T (V_y^*)^{-1} (Y_{0,obs}(s_i, \cdot) - Y_{0,pred}^m(s_i, \cdot))$, where $Y_{0,pred}(s_i, \cdot)$ and V_y^* are computed using Rao-Blackwellization (Gelfand and Smith, 1990) as follows, $Y_{0,pred}(s_i, \cdot) = \frac{1}{m} \sum Y_{0,pred}^m(s_i, \cdot)$ and $V_y^* = \frac{1}{m} [V_y^m + Y_{0,pred}^m(s_i, \cdot) Y_{0,pred}^m(s_i, \cdot)^T] -$

$Y_{0,pred}(s_i, \cdot)Y_{0,pred}(s_i, \cdot)^T$. For this region, we again use a chi-squared distribution with T degrees of freedom to create a cutoff; however, here, this will only be approximate, with the approximation improving as the posterior of (β, η, τ^2) concentrates around the true values.

If the original data were transformed, the predictive variance will be needed to obtain the posterior predictive mean on the original scale. For example, in the application, we log transformed the responses. So, the prediction on the original scale, X_{pred} , is given as

$$X_{pred} = \exp\{Y_{pred} + V_y/2\}.$$

We will compute MSPE on both the original scale and the transformed scale of the data .

In the following example, we will examine MSPE and coverage overall, and for each process separately.

3 Data Example

We illustrate this methodology on pollutant data from the LA area. There were a total of 35 sites with pollutant monitors, 5 of which monitored both particulate matter (PM) and ozone and 30 which only monitored ozone; a map of the sites is given in Figure 1. We analyze data collected during a 54 day period during the summer in 1995, July 10-Sept. 1 (this period was chosen since it contained no missing data); however, we point out that that the time window chosen was chosen for ease of demonstration of these models. These models easily accommodate intermittent missing data (even if missing by design) as discussed in Section 2.5. We also mention that we undertake a simple analysis of this data here. A more definitive analysis would include many extensions which we discuss more thoroughly in Sections 3.7 and 4. One objective of this modeling was to use ozone to help predict PM given the sparsity of the PM sites.

A log transformation of both ozone and PM seemed to provide a good approximation to normality and constant variance. This agrees with previous work by Carroll et al. (1997) on ozone and Cressie et al. (1999) on PM. The design matrix ultimately chosen, W , consisted

of 9×1 design vectors which included a set of ozone-specific components: an intercept, terms for both latitude and longitude and an interaction, and a quadratic in time; and a set of PM-specific components: an intercept and a quadratic in time. We also considered some additional design vectors that included/excluded spatial components for the mean structure for PM/ozone. Some weather covariates were also available, but there was so much missingness here that they could not be used without explicitly building models for them. Ultimately, the most complex mean structure 'fit' the data best (by DIC), but resulted in poorer cross-validation predictive accuracy.

A summary of the data by site appears in Table 1.

3.1 Neighborhood Structure

We consider the following possible neighbors of $Y_k(s_{ki}, t)$

1. $Y_k(s_{ki}, t + 1)$. The pollutant value at the same location on the subsequent day.
2. $Y_k(s_{ki}, t - 1)$. The pollutant value at the same location on the previous day.
3. $\{Y_k(s_{kl}, t) : l \neq i\}$. The pollutant values on the same day at other locations.
4. $\{Y_j(s_{ji}, t) : j \neq k\}$. The pollutant values of other pollutants on the same day at the same location.

Let $c_{(i,i'),(t,t')}^{(k,k')}$ denote the element of the $C(\eta)$ matrix corresponding to the pair of values $(Y_k(s_{ki}, t), Y_{k'}(s_{k'i'}, t'))$. Using the neighborhood structure given above, we specify the $C(\eta)$ matrix to have the following structure,

1. $c_{(i,i),(t,t+1)}^{(k,k)} = \eta_t^{(k)}$
2. $c_{(i,i),(t,t-1)}^{(k,k)} = \eta_t^{(k)}$
3. $c_{(i,l),(t,t)}^{(k,k)} = \eta_S^{(k)} f(s_i, s_l)$
4. $c_{(i,i),(t,t)}^{(k,j)} = \eta_S^{(k,j)}$

and 0 otherwise. In this application, we set $f(s_i, s_l) = \min_{i', l'}(d_{i', l'})/d_{i, l}$ where $d_{i, l}$ is the Euclidean distance between the two sites. This is similar to the choice of Cressie and Chan (1989), though they set the dependence to 0 when the distance between locations exceeded some data determined threshold. For a good approach for determining distance based neighborhood structures using geostatistical techniques, we refer the reader to Hrafnkellson and Cressie (2003), which we will discuss further in Section 4. For block source data, $f(s_i, s_l)$ might be chosen to be an indicator of whether sites s_i and s_l are neighbors or some measure of the 'quality' of the neighbors (see Section 4.4).

3.2 Choice of the conditional covariance matrix, M^*

For the M^* matrix, we consider two specifications. For the first, we assume a common conditional variance over space, i.e., $M^* = \tau^2 I_p$. For the second, we assume a separate conditional variance for each process (pollutant); in our application, τ_1^2 and τ_2^2 . The elements of the $C(\eta)$ matrix then need to be adjusted for the matrix $M^{-1}(I_{pT} - C(\eta))$ to be symmetric. The elements of the $C(\eta)$ matrix will remain as specified above except for the elements of the $C(\eta)$ matrix corresponding to cross-pollutant neighbors. These will have the following structure,

1. $c_{(i,i),(t,t)}^{(k,j)} = \eta_S^{(j,k)} \tau_j / \tau_k$
2. $c_{(i,i),(t,t)}^{(j,k)} = \eta_S^{(j,k)} \tau_k / \tau_j$

where the first term is in the upper triangle and the second term is in the lower triangle of the $C(\eta)$ matrix.

3.3 Posterior sampling details

The full conditional for τ^2 in the model with $M^* = \tau^2 I_p$ will be inverse gamma. For the more general specification of M^* , the $C(\eta)$ matrix will contain the diagonal elements of the M^* matrix. For this choice of M^* , we use a random walk Metropolis-Hastings algorithm. For the neighborhood structure, given in Section 3.1, we can compute 'necessary' intervals for

positive definiteness (i.e., values outside the interval, correspond to a non-positive definite matrix). For $M = \tau^2 I$ and the specific M^* described in Section 3.2, necessary intervals, given our neighborhood structure, for $\eta_s^{(k)}$ are $(1/\psi_{k,max}, 1/\psi_{k,min})$, where $\psi_{k,min}$ and $\psi_{k,max}$ are the smallest and largest eigenvalues respectively of the distance matrices for each pollutant (k), i.e., the $N_k \times N_k$ matrix of pairwise distances (as specified in Section 3.1 with 0's on the main diagonal). In addition, the following bound must be satisfied which provides additional restrictions on $\eta_t^{(k)}$ given $\eta_s^{(k)}$ and vice versa, $|\eta_t^{(k)}| < \frac{1}{\max_l \lambda_l} (1 - \eta_s^{(k)} \psi_{k,max})$. Here, we have assumed that $\eta_s^{(j,k)} = 0$; the intervals are further restricted if this is not the case. In Section 3.4, we use these bounds to help interpret the magnitude of the estimated dependence parameters (η). However, we cannot use a logit transformation based on this last interval as the bounds for $\eta_t^{(k)}$ depend on the current value of $\eta_s^{(k)}$ and vice versa. However, we can use the first set of bounds for $\eta_s^{(k)}$ and simpler bounds on $\eta_t^{(k)}$, $(-.5, .5)$, based on the eigenvalues of the ρ matrix, λ .

The Gibbs sampling algorithm converged quickly using several starting values. 50,000 iterations were run, with the first 1000 thrown out as burn-in; computing time for complex models like Models IV and V was four hours on a PC with a pentium 4, 1.8 GHz processor. For Model V, we broke the dependence parameters, η 's into two blocks: 1) $(\eta_s^{(1)}, \eta_s^{(2)}, \eta_t^{(1)}, \eta_t^{(2)})$, and 2) $\eta_s^{(1,2)}$. The dependence parameters mixed quite well using the Hybrid MC algorithm though the cross-pollutant parameter, $\eta_s^{(1,2)}$ mixed more slowly than the others. Convergence was assessed by examining trace plots from four chains.

3.4 Model fit and inference on η parameters

We fit five dependence models as given in Table 2. Model fit, as discussed in Section 2.5, was assessed via the DIC. DIC values for the given models, are given in Table 3. Clearly, the model with all the spatial and temporal dependence parameters fit best.

The estimates of η , the dependence parameters, and τ_k^2 , $k = 1, 2$, the variance parameters, for the best fitting model, Model V, appear in Table 4. Using the information on the necessary

intervals/bounds in Section 3.3, all of the spatial and temporal dependence parameters are quite close to their upper bounds. As the intervals in Section 3.3 do not include $\eta_s^{(1,2)}$, we did some simulations conditional on the other η parameters and found this was also in fact near its upper bounds as well. So, clearly, both the spatial and temporal dependence is strong in this data.

3.5 Prediction Error

For prediction, for each sub-site, 5100 iterations were run, with a burn-in of 100; twenty-two hours was required to run the prediction for all the sub-sites in the model. Predictive ability as measured by MSPE is given in Tables 5-8 for prediction on the log and original scale both spatially and temporally. We use the posterior means of τ_1^2 and τ_2^2 and Table 1 to help calibrate the magnitude of the prediction errors for PM and ozone respectively, below.

The overall gains in spatial prediction when comparing the independence model to the best dependence model, Model V, was to decrease MSPE by about 30%. The magnitude of the prediction error was about twice the value of the respective posterior means of τ^2 (cf: Table 4). To assess the gains in PM prediction by modelling it jointly with ozone, we compare models IV and V. The gains in predicting PM here are not large, on the order of a 2-4% reduction in MSPE. The bias in general was quite small (cf. Table 1).

The superior temporal prediction for PM versus ozone is mainly due to the fact that an AR(1) type model fit PM very well, but not ozone (not shown). However, we do see a decrease in MSPE of over 30% for ozone when including temporal dependence in the model. The huge differences between the MSPE for PM was due to very poor prediction for the independence model on the first day. We can see that in Model V, the magnitude of the prediction error for PM was less than the posterior mean for τ_1^2 while it was about twice the value of τ_2^2 for ozone. This indicates solid temporal prediction for PM.

3.6 Coverage results

To obtain some absolute measure of model fit, we examined the coverage probabilities (based on 95% confidence regions) for Model V using the cross-validation approach. For the PM spatial predictive coverage, we observed coverages between .94 and 1.0 for four of the five sites, with one site having coverage of 0.0; for ozone, we observed coverage between .94 and 1.0 for 28 of the 35 sites, with 7 sites having less than .04 coverage. Thus, for both PM and ozone, the model appears to be fitting about 80% of the sub-sites well. This agreed with the marginal credible regions which covered 80% of the sub-sites. For PM temporal predictive coverage, coverage was over 90% for both days while for ozone, coverage was about .98 for the last day, but .00 for the first day.

We did some further exploration of the sub-sites with poor coverage. These tended to be sites with the fewest 'quality' neighbors (for a discussion of 'quality' neighbors, see Section 4.4) and with considerable variability over time (relative to the other sub-sites; see Table 1). Poorer prediction at sub-sites with not too many 'quality' neighbors is not surprising.

3.7 Model Expansions

These results suggest further model expansion from several perspectives. In terms of spatial dependence, nonstationary spatial structures should be considered, especially given the topography in this region of the United States, as well as fine tuning the distance based neighborhood structure using the approach of Hrafnkellson and Cressie (2003) (see Section 4.5). Other extensions might include adjusting the M matrix to account for the 'quality' of neighbors (see Section 4.4) and taking into account the measurement error in the monitors (Section 4.2). We will discuss some details about these expansions and others in the next section.

4 Model Generalizations and Extensions

In the following, we will discuss several generalizations and extensions to the model, both with respect to the neighborhood structure and incorporating additional sources of variability.

4.1 More than lag 1 temporal dependence

Only lag 1 temporal dependence has been considered here. Increasing the temporal neighborhood structure beyond the single nearest neighbor creates major difficulties with computations. Such matrices will have a block Toeplitz form for which the inverse can be computed efficiently (Dietrich, 1991), but checking positive definiteness and computing the determinant are not computationally tractable.

4.2 Measurement Error

Another extension would be to build a hierarchical (mixture) model to allow for measurement error. These models are most appropriate when the variance of the measurement error is known or there are replications. In our setting, we did not know the measurement error variances and we did not have replications. However, the measurement error variance still could be identified by assuming independence of the measurement errors, conditional on the true values of the process; the measurement error variance estimated in this way would likely be confounded by residual error not explained in the spatial dependence model for the true process (see below). An alternative approach would be to put an informative prior on σ^2 .

The extension to our model is as follows,

$$Y \sim N(\theta, D) \tag{9}$$

$$\theta \sim N(W\beta, [M^{-1}(\tau^2)(I_{pT} - C(\eta))]^{-1}) \tag{10}$$

with D a diagonal matrix often specified as $\sigma^2 I$. We have added a measurement error component into the model and the true process of interest is now represented by θ , not

y . How does this impact computing? Sampling from the full conditional distributions of β, η, τ^2 will be the same, but with Y replaced by θ . The full conditional for σ^2 will be inverse gamma (assuming a conjugate prior). The full conditional distribution of θ will be normal and the tools used in this paper to deal with $M^{-1}(I_{pT} - C(\eta))$ can be used for computing this distribution. Finally, intermittent missing data can easily be handled more easily within this model as the missing Y 's can be sampled from a $N(\theta, D)$ distribution.

For count or binary data, (9) could be replaced by a Poisson or logistic regression model with θ now being a random effect.

4.3 Extension to flexible MCAR models

Gelfand and Vounatsou (2003) consider linear transformations of the data/random effects to extend the MCAR models of Mardia (1988) to allow a separate spatial parameter for each of the multivariate processes. We will give some details on incorporating these ideas into our model. Consider a $p \times p$ linear transformation matrix G , which is block diagonal with blocks of size K . We can modify the model proposed in this paper by replacing the covariance matrix $M^{-1}(I_{pT} - C(\eta))$ with the matrix $(G \otimes I_T)M^{-1}(I_{pT} - C(\eta))(G \otimes I_T)'$. The benefit of extending the model using this approach is a more flexible specification for multivariate spatial dependence with simple constraints on the spatial dependence parameters for positive definiteness; thus, no check would need to be made on the eigenvalues, γ_i , as in Theorem I. In addition, Theorem I will hold as before with appropriate modification to the matrix B . Care, however, must be taken to account for the impact of G on the temporal neighborhood matrix, V_t .

4.4 Adjustment for 'quality' of neighbors

In the air pollution dataset here, and in many other settings with an irregular lattice, there is an asymmetry in terms of the number of neighbors (for block sources) or the number and quality of neighbors (with point sources). The 'adjustment' often made to the conditional variance in CAR models with block sources is to divide a constant unknown variance by

the number of neighbors, τ^2/n_i (Sun et al, 1999). The motivation for such an adjustment is that we would expect the conditional variance of sites with more neighbors to be smaller than sites with fewer neighbors. Adjustments are also frequently made for the quality of neighbors as well; e.g., how much of the total boundary do they share. In point source data, often all sites are neighbors to each other. We might construct a similar adjustment to the standard CAR models in this case by dividing τ^2 by a measure of the 'quality' of the neighbors. We might define this quantity, n^* , by the inverse of the harmonic average of the distances from the current site to all other sites. Since the neighborhood structure in this model is assumed static over time, n^* will be constant over time and computations will proceed under the $M = I_T \otimes M^*$ case. Appropriate adjustments then need to be made to the conditional mean (through the C matrix) to ensure $M^{-1}(I_{pT} - C(\eta))$ remains symmetric. This idea will be explored in future work and might improve the prediction for the pollution example as discussed in Section 3.6.

4.5 Alternative neighborhood structures and nonstationarity

The temporal neighborhood structure used in the example could be supplemented by including the following additional neighbors for $Y_k(s_{ki}, t)$; $Y_k(s_{kl}, t + 1), l \neq i$ (lag 1 temporal correlation across sites) and/or $Y_j(s_{jl}, t + 1), l \neq i, j \neq k$ (lag 1 temporal correlation across sites and pollutant). In addition, to better calibrate distance-based neighbors, the techniques in Hrafnkellson and Cressie (2003) could be implemented which involves, among other things, finding a range of dependence for the neighbors by 'equating' the neighborhood structure to the results from fitting variograms using geostatistical techniques. In this paper, the distance-based neighborhood structure was specified without using such techniques as in Kaiser et al. (2002) among others.

The neighborhood structure proposed for our models might be characterized as stationary. To extend to a nonstationary dependence, we can allow the neighborhood dependence parameters only to operate locally and/or allow them to vary over space. The deformation of

distance approach (Sampson and Guttorp, 1992) might be a reasonable methodology to try to integrate into this model. Care must be taken in order to ensure the conditional variance matrix M^* and the neighborhood matrix, $C(\eta)$ are specified appropriately so as to make $M^{-1}(I_{pT} - C(\eta))$, symmetric.

5 Discussion

We have proposed a conditionally specified space-time model which provides both tractable computations and flexibility in specifying the spatial dependence structure while allowing for simple temporal dependence. This provides another class of models that can be considered for space-time data with either point or block sources.

The multi-pollutant example illustrated the potential gains from modeling multiple space-time processes jointly via the slightly increased predictive accuracy and greatly improved model fit (DIC). Larger gains would likely have been seen in the example by modeling pollutants that are more related than PM and ozone and by introducing some of the model expansions discussion in Sections 3.7 and 4.

In addition to 'modeling' the pollutant field as an end in and of itself, the 'output' from these models has other uses. For example, these multi-pollutant space-time models might be used as input to health effects models. However, the health effects data are usually at a different level of aggregation than the pollutant data. In particular, the pollutant data are point sources and the health effects data tend to be over areas. So, in some sense, an average exposure (pollutant level) for the area would be desired. This could be accomplished by creating a fine grid over the areas, constructing the appropriate $C(\eta)$ matrix on this grid, predicting the pollutants at each point on the grid (Hrafnkelsson and Cressie, 2003), and then averaging. Such an approach would avoid the change of support problem discussed recently in Gelfand et al. (2001). In addition, the posterior draws from the Gibbs sampler can be used to address the uncertainty of the pollutant exposure in the health effects models. Recent work on such modeling has been done by Van den Eeden (2001), Zhu et al. (2003),

and Dukic et al. (unpublished manuscript). For block sources, we might use these models to fill in spatial covariates for a regression model where some of the covariates (exposures) are missing in certain blocks.

In the LA data, weather variables were also collected. However, there was considerable missingness. These might be used in both the spatial model for the pollutants and the health effects models by fitting additional spatial models to the weather data to obtain a predicted weather field that could be used as covariates in both the pollutant and health effects models.

A criticism of conditionally specified models is that the marginal distribution of the original sites changes when we add in additional sites; i.e., lack of marginal invariance. For block sources this is less of an issue than with point sources. However, the ultimate goal of the modeling may either be a good fit of the model to the data and/or good prediction; the model should be judged more on these criteria than the lack of marginal invariance property.

An alternative approach for prediction (cf: Section 2.5) would be to augment the $C(\eta)$ matrix with the site(s) to be predicted, Y_0 , and then to sample from the posterior distribution with the additional site(s) in the model, similar to the data augmentation approach used to 'fill-in' intermittent missing data. This approach may potentially result in slightly different predictions and posterior distributions of the parameters since the marginal distribution of the observed sites is altered when adding in the site to be predicted. This was the approach used in Hrafnkelsson and Cressie (2003).

The fit of the covariance structure could be assessed via posterior predictive checks (Gelman et al., 2003). Such checks have been used recently to assess the fit of the correlation structure in complex hierarchical models (see, e.g., Ilk and Daniels, 2005).

Future work will include doing many of the model expansions discussed here for a more definitive analysis of the pollutant data and making comparisons of such models with geostatistics models with a separable space-time structure. In addition, we will report on the propriety of the posterior under improper priors on the regression coefficients, β , and the diagonal matrix of variances, M , including derivation of default reference priors for the

dependence parameters. We will also explore efficient sampling algorithms as the neighborhood/dependence structure gets more complex.

Acknowledgments

The authors would like to thank Dr. Stephen van den Eeden for providing the 1995 air pollution data, Dr. Mark Kaiser for helpful discussions on conditionally specified models, and three referees whose comments greatly improved the manuscript. Part of this work was funded by NSF DMS-9816630 and by NIH CA85295.

REFERENCES

- Anderson, T. W. (1984), An introduction to multivariate statistical analysis, John Wiley & Sons (New York; Chichester).
- Berger, J.O., De Oliveira, V., and Sanso, B. (2001) Objective Bayesian analysis of spatially correlated data. *Journal of the American Statistical Association*, 96, 1361-1374.
- Besag, J.E. (1974) Spatial interaction and the statistical analysis of lattice systems (with discussion). *Journal of the Royal Statistical Society B*, 36, 192-225.
- Best, Nicola G. , Arnold, Richard A. , Thomas, Andrew , Waller, Lance A. , and Conlon, Erin M. (1999), “Bayesian models for spatially correlated disease and exposure data”, In *Bayesian Statistics 6*, Oxford: Oxford University Press, 131-156
- Brown, P.E., Diggle, P.J., Lord, M.E., Young, P.C. (2001) Space-time calibration of radar rainfall data. *Applied Statistics*, 50, 221-241.
- Carlin, B.P., and Banerjee, S. (2002) Hierarchical Multivariate CAR models for spatially correlated survival data. In *Bayesian Statistics 7*. Oxford: Oxford University Press, 45-64.
- Carroll, R.J., Chen, R., George, E.I., Li, T.H., Newton, H.J., Schmiediche, H., Wang, N. (1997) Ozone exposure and population density in Harris County, Texas (with discussion). *Journal of the American Statistical Association*, 92, 392-415.
- Conlon, E. and Waller L. (1998) Flexible neighborhood structures in hierarchical models for disease mapping. *University of Minnesota Biostatistics Research Report*, No. rr98-018.

- Cressie, N.A.C., Chan, N.H. (1989) Spatial modeling of regional variables, *Journal of the American Statistical Association*, 84, 393-401.
- Cressie, N., Kaiser, M.S., Daniels, M.J., Aldworth, J., Lee, J., Lahiri, S.N., Cox, L.H. (1999) Spatial analysis of particulate matter in an urban environment. In *geoENV II - Geostatistics for Environmental Applications*, Gomez-Hernandez, J, Soares, A., Froidevaux, R. (eds). Kluwer: Dordrecht, 41-52 .
- Dietrich, C.R. (1991) Computationally efficient cholesky factorization of a covariance matrix with block toeplitz structure. *Journal of Statistical Computation and Simulation*, 45, 210-218.
- Gelfand, A.E. and Smith, A.F.M. (1990) Sampling based approaches to calculating marginal densities. *Journal of the American Statistical Association*, 85, 398-409.
- Gelfand, A.E., Vounatsou, P. (2003) Proper multivariate conditional autoregressive models for spatial data analysis. *Biostatistics*, 4, 11-25.
- Gelfand, A.E., Zhu, L., Carlin, B.P. (2001) On the change of support problem for spatio-temporal data, *Biostatistics*, 2, 31-45.
- Gelman, A., Carlin, J.B, Stern, H., Rubin, D. (2003) Bayesian data analysis. 2nd edition. Chapman and Hall/CRC.
- Handcock, M.S., and Wallis, J.R. (1994) An approach to statistical spatial-temporal modeling of meteorological fields. *Journal of the American Statistical Association*, 89, 368-378.
- Haas, T.C. (1995) Local prediction of a spatio-temporal process with an application to wet sulfate deposition. *Journal of the American Statistical Association*, 90, 1189-1199.
- Haas, T.C. (2002) New systems for estimating, and predicting a multivariate spatio-temporal process, *EnvironMetrics*, 13, 311-332.
- Hrafinkelsson, B., Cressie, N. (2003) Hierarchical modeling of count data with application to nuclear fall-out. *Environmental and Ecological Statistics*, 10, 179-200.
- Ilk, O., Daniels, M.J. (2005) Marginalized transition random effects models for multivariate longitudinal binary data. Submitted.

- Kaiser, M.S., Daniels, M.J., Furakawa, K., Dixon, P. (2002) Analysis of particulate matter air pollution using Markov random field models of spatial dependence. *Environmetrics*, 13, 615-628.
- Kaiser, M.S. and Cressie, N. (2000) The construction of multivariate distributions from Markov random fields. *Journal of Multivariate Analysis*, 73, 199-220.
- Kibria, B., Golam, M., Sun, L., Zidek, J.V., Le, N.D. (2002) Bayesian spatial prediction of random space-time fields with application to mapping PM2.5 exposure. *Journal of the American Statistical Association*, 97: 112-124.
- Kim, Hoon, Sun, Dongchu, and Tsutakawa, Robert K (2001), A bivariate Bayes method for improving the estimates of mortality rates with a twofold conditional autoregressive model, *Journal of the American Statistical Association*, 96 (456), 1506-1521
- Knorr-Held, L. (2000) Bayesian modeling of inseparable space-time variation in disease risk. *Statistics in Medicine*, 19, 2555-2567.
- Knorr-Held, L., Rue, H. (2002) On Block updating of Markov Random Field Models for Disease Mapping. *Scandinavian Journal of Statistics*, 29, 597-614.
- Le, N.D., Sun, W., Zidek, J.V. (1997) Bayesian multivariate spatial interpolation with data missing by design. 59, 501-510.
- Mardia, K.V. (1988) Multi-dimensional multivariate Gaussian Markov random fields with application to image processing. *Journal of Multivariate Analysis*, 24, 265-284.
- Mardia, K.V., Goodall, C., Redfern, E.J., Alonso, F.J. (1998) The Kriged Kalman Filter (with discussion). *TEST*, 7, 217-285.
- Neal, R.M. (1994) An improved acceptance procedure for the hybrid monte-carlo algorithm. *Journal of Computational Physics*, 111, 194-203.
- Niu, X-F, McKeague, I.W., Elsner, J.B. (2003) Seasonal space-time models for climate systems. *Statistical Inference for Stochastic Processes*, 6, 111-133.
- Sampson, P.D. and Guttorp, P. (1992) Nonparametric estimation of nonstationary spatial covariance structure, *Journal of the American Statistical Association*, 87, 108-119
- Shaddick, G., Wakefield, J. (2002) Modelling daily multivariate pollutant data at multiple sites. *Applied Statistics*, 51, 351-372.

- Spiegelhalter, D.J., Best, N.G., Carlin, B.P. and van der Linde, A. (2002) Bayesian measures of model complexity and fit (Pkg: p583-639) *Journal of the Royal Statistical Society, Series B*, 64, 583-616.
- Stein, M.L. (2003) Space time covariance functions. *The University of Chicago, Center for Integrating Statistical and Environmental Science* Technical Report, No. 4.
- Stroud, J.R., Muller, P., Sanso, B. (2001) Dynamic models for spatiotemporal data. *Journal of the Royal Statistical Society, Serie B*, 63, 673-689.
- Sun, Dongchu , Tsutakawa, Robert K. , and Speckman, Paul L. (1999), Posterior distribution of hierarchical models using CAR(1) distributions, *Biometrika*, 86 , 341-350
- Van den Eeden, S. (2001) Air pollution and morbidity and mortality among California Kaiser Permanente members. Kaiser Permanente Research Report.
- Wikle, C.K., Berliner, L.M., Cressie, N. (1998) Hierarchical Bayesian space-time models. *Environmental and Ecological Statistics*, 5, 117-154.
- Wikle, C.K. and Cressie, N. (1999) A dimension-reduced approach to space-time Kalman filtering. *Biometrika*, 86, 815-829.
- Zhu, L, Gelfand, A.E., and Carlin, B.P. (2003) Hierarchical regression with misaligned spatial data: Relating ambient ozone and pediatric asthma ER visits in Atlanta. *Environmetrics*, 14, 537-557.
- Zidek, J.V., Sun, W., Le, N.D. (2001) Designing and integrating composite networks for monitoring multivariate Gaussian pollution fields. *Applied Statistics*, 49, 63-79.
- Zidek, J.V., Sun, W, Le, N.D., Ozkaynak, H. (2002) Contending with space-time interaction in the spatial prediction of pollutions: Vancouver's hourly ambient PM₁₀ field. *Environmetrics*, 13, 595-613.




Figure 1: PM and ozone monitor locations.

Pollutant	Site No.	Mean	Variance	Min	Max
PM	1	3.93	.077	3.42	5.27
	2	3.70	.056	3.13	4.76
	3	4.38	.059	3.75	4.81
	4	4.21	.112	3.36	5.35
	5	3.79	.040	3.47	4.34
ozone	1	3.22	.042	2.73	3.61
	2	3.15	.073	2.65	3.65
	3	3.78	.059	3.35	4.25
	4	3.68	.082	3.21	4.27
	5	3.38	.082	2.73	3.85
	6	2.67	.094	1.72	3.24
	7	3.66	.046	3.30	4.10
	8	3.73	.121	2.83	4.55
	9	3.46	.051	2.96	4.05
	10	3.57	.118	3.14	4.60
	11	4.00	.282	1.97	4.59
	12	3.28	.089	2.48	3.88
	13	3.13	.092	2.45	3.71
	14	3.83	.061	3.41	4.38
	15	3.30	.044	2.59	3.71
	16	3.79	.043	3.23	4.19
	17	3.96	.054	3.39	4.45
	18	3.24	.067	2.64	3.73
	19	3.81	.058	3.17	4.32
	20	3.92	.075	3.36	4.35
	21	3.63	.095	3.05	4.28
	22	3.66	.053	3.29	4.11
	23	3.54	.047	3.04	4.01
	24	3.86	.058	3.34	4.41
	25	3.43	.038	3.02	3.84
	26	3.53	.040	3.13	3.96
	27	3.53	.068	3.08	4.03
	28	4.04	.083	3.37	4.54
	29	3.48	.062	2.82	4.04
	30	3.95	.065	3.53	4.41
	31	3.80	.062	3.35	4.60
	32	3.59	.051	3.03	4.60
	33	3.75	.074	3.33	4.28
	34	3.84	.139	2.57	4.60
	35	3.31	.046	2.62	3.73

Table 1: Summary statistics for each site over time on the log scale. Sites 1-5 are measured for both PM and ozone.

Model	Parameters	
	Spatial	Temporal
I	-	-
II	-	$(\eta_t^{(1)}, \eta_t^{(2)})$
III	$(\eta_s^{(1)}, \eta_s^{(2)})$	-
IV	$(\eta_s^{(1)}, \eta_s^{(2)})$	$(\eta_t^{(1)}, \eta_t^{(2)})$
V	$(\eta_s^{(1)}, \eta_s^{(2)}, \eta_s^{(1,2)})$	$(\eta_t^{(1)}, \eta_t^{(2)})$

Table 2: Models Fit

Model	DIC	p_D
I	1259	11.7
II	269	13.8
III	755	14.7
IV	99.7	14.9
V	78.1	14.3

Table 3: DIC and effective number of parameters (p_D) for models fit

Parameter	Posterior mean (CI)
$\eta_{space}^{(1)}$.046 (.005,.099)
$\eta_{space}^{(2)}$.043 (.035,.051)
$\eta_{space}^{(1,2)}$.065 (.033,.094)
$\eta_{time}^{(1)}$.42 (.37,.45)
$\eta_{time}^{(2)}$.39 (.37,.41)
τ_1^2	.055 (.046,.066)
τ_2^2	.049 (.045,.052)

Table 4: Posterior means and 95% credible intervals for the spatial dependence parameters. (1) corresponds to PM and (2) to ozone.

Model	Overall		PM		Ozone	
	MSPE	Bias	MSPE	Bias	MSPE	Bias
I	.119	.003	.169	.000	.112	.003
IV	.094	.004	.107	-.003	.092	.005
V	.093	.005	.103	-.004	.091	.007

Table 5: MSPE and bias overall and broken down by pollutant for predicting one site at all times on log scale

Model	Overall		PM		Ozone	
	MSPE	Bias	MSPE	Bias	MSPE	Bias
I	239	0.4	735	.15	168	.41
IV	180	0.0	508	-.6	133	.1
V	178	0.0	498	-.9	133	0.1

Table 6: MSPE and bias overall and broken down by pollutant for predicting one site at all times on original scale

Model	Overall		PM		Ozone	
	MSPE	Bias	MSPE	Bias	MSPE	Bias
I	1.26	-.30	9.1	-2.1	.133	-.04
V	.092	-.09	.030	-.08	.100	-.10

Table 7: MSPE and bias overall and broken down by pollutant for predicting all sites at the first and last times on log scale

Model	Overall		PM		Ozone	
	MSPE	Bias	MSPE	Bias	MSPE	Bias
I	547	-5.5	2730	-35	235	-1.3
V	190	-4.0	87	-4.0	204	-3.9

Table 8: MSPE and bias overall and broken down by pollutant for predicting all sites at the first and last times on original scale