Published in final edited form as:

J R Stat Soc Ser C Appl Stat. 2020 June ; 69(3): 681–696. doi:10.1111/rssc.12407.

A spatially varying distributed lag model with application to an air pollution and term low birth weight study

Joshua L. Warren,

Yale University, New Haven, USA

Thomas J. Luben,

US Environmental Protection Agency, Research Triangle Park, USA

Howard H. Chang

Emory University, Atlanta, USA

Summary.

Distributed lag models have been used to identify critical pregnancy periods of exposure (i.e. critical exposure windows) to air pollution in studies of pregnancy outcomes. However, much of the previous work in this area has ignored the possibility of spatial variability in the lagged health effect parameters that may result from exposure characteristics and/or residual confounding. We develop a spatially varying Gaussian process model for critical windows called ‘SpGPCW’ and use it to investigate geographic variability in the association between term low birth weight and average weekly concentrations of ozone and PM_{2.5} during pregnancy by using birth records from North Carolina. SpGPCW is designed to accommodate areal level spatial correlation between lagged health effect parameters and temporal smoothness in risk estimation across pregnancy. Through simulation and a real data application, we show that the consequences of ignoring spatial variability in the lagged health effect parameters include less reliable inference for the parameters and diminished ability to identify true critical window sets, and we investigate the use of existing Bayesian model comparison techniques as tools for determining the presence of spatial variability. We find that exposure to PM_{2.5} is associated with elevated term low birth weight risk in selected weeks and counties and that ignoring spatial variability results in null associations during these periods. An R package (SpGPCW) has been developed to implement the new method.

Keywords

Air pollution; Critical windows; Gaussian process; Spatially varying coefficients; Term low birth weight

Address for correspondence: Joshua L. Warren, Department of Biostatistics, Yale University, 60 College Street, New Haven, CT 06510, USA. joshua.warren@yale.edu.

Supporting information

Additional ‘supporting information’ may be found in the on-line version of this article: ‘On-line supplementary material’.

1. Introduction

Distributed lag models (DLMs) have been widely used in environmental epidemiology to quantify associations between environmental exposures occurring at different lagged periods and adverse health outcomes (Samet *et al.*, 2000; Schwartz, 2000; Zanobetti *et al.*, 2000; Peel *et al.*, 2005; Welty and Zeger, 2005; Dominici *et al.*, 2006; Zanobetti and Schwartz, 2009). DLMs have also been used to identify and estimate critical exposure windows: specific exposure periods during which an individual may be especially sensitive to environmental contaminants (Hu *et al.*, 2006; Darrow *et al.*, 2010; Sanchez *et al.*, 2010; Leon Hsu *et al.*, 2015). Uncovering these critical windows may provide an improved understanding of possible mechanisms by which air pollution impacts health, and this was recently identified as a strategic goal by the National Institute of Environmental Health Sciences (National Institute of Environmental Health Sciences, 2012). In addition, identifying and characterizing critical exposure windows will enable enhanced guidance and aid in protecting sensitive populations.

The use of DLMs for identifying critical exposure windows has led to the development of statistical methodology to overcome remaining analytic challenges that often arise from different areas of application. Driven by applications in reproductive epidemiology (i.e. the effect of pollution on preterm birth), Warren *et al.* (2012a) used Gaussian processes to model exposure to average weekly pollution concentrations across pregnancy while also allowing for spatial variability in the lagged health effect parameters and an extension to two pollutants, whereas Warren *et al.* (2012b) extended this work to the semiparametric Bayesian setting and for a multivariate outcome (i.e. cardiac congenital anomalies). Working in the survival statistics setting (i.e. treating preterm birth as a time-to-event outcome), Chang *et al.* (2015) developed dynamic Gaussian processes to estimate weekly risk parameters that varied by exposure and health outcome time periods. In a study of children's health, Wilson *et al.* (2017) introduced a method to explore the possibility that critical exposure windows vary among subgroups of the study population. To analyse the effect of mixtures in the DLM framework, Bello *et al.* (2017) proposed the use of lagged weighted quantile sums whereas Liu *et al.* (2017) extended the Bayesian kernel machine regression framework (Bobb *et al.*, 2014) to handle multiple exposure time periods. Wilson *et al.* (2019) developed a Bayesian kernel machine regression DLM for the analysis of pollution mixtures based on the use of time-weighted exposures in the specified kernel function. Recently, Warren *et al.* (2019) considered the use of smoothed Bayesian variable selection for a more explicit definition of a critical window. We also note that new DLM methodology that is not directly related to critical window estimation has been developed, including for estimating interactions (Heaton and Peng, 2013; Chen *et al.*, 2019) and non-linear associations (Gasparrini *et al.*, 2010).

Although various advances have been made, currently few studies have investigated the possibility of spatial variation in the lagged regression parameters. Contributors to such spatial variation can include exposure characteristics (e.g. spatially varying pollutant composition and spatial differences in exposure measurement error) and/or characteristics of the study population that may be omitted from the analysis or are included but represent imperfect proxies of important covariates or confounders (e.g. race or ethnicity and education level). To account for spatial variation, Warren *et al.* (2012a, b) used geostatistical

methods to introduce spatial correlation between lagged regression parameters corresponding to different discrete spatial regions while defining correlation as a function of distance between region centroids. However, this distance may be inappropriate when the spatial regions are oddly shaped and/or their sizes vary greatly (Banerjee *et al.*, 2014). Additionally, introducing a model using distance-based spatial correlation can result in increased computational complexity as the number of unique locations grows because of the reliance on a Gaussian process with spatial covariance function and the need to invert a large matrix during each iterative step of the model fitting algorithm (typically carried out by using Markov chain Monte Carlo sampling techniques) (Heaton *et al.*, 2018).

Outside the search for critical windows, but working with DLMs, Baek *et al.* (2016) allowed for spatial region-specific deviations from a global lag structure while using smoothing splines to restrict the behaviour of the lagged parameters. However, spatial correlation between the lagged parameters was not directly considered. Instead, the location-specific lagged parameters were independent across space, though correlated with each other in the case of the intercept and slope parameters. A spatially structured random effect was included outside the model for the lagged parameters to account for spatial correlation between responses.

In this work, we offer an alternative formulation to the distance-based spatial correlation Gaussian process DLM framework by allowing for areal level spatial correlation between lagged parameters by using a proper version of the multivariate conditional auto-regressive (MCAR) model (Leroux *et al.*, 2000; Kavanagh *et al.*, 2016) while incorporating regularization across discrete exposure time via an exponential correlation function. We apply a spatially varying Gaussian process model for critical windows called ‘SpGPCW’ to investigate spatial variability in the association between term low birth weight and exposure to weekly ambient ozone and particulate matter less than 2.5 μm in diameter, $\text{PM}_{2.5}$, across pregnancy. To do this, we analyse birth certificate records from all births between 2005 and 2008 in the Raleigh–Durham–Chapel Hill combined statistical area (RDCHCSA) of North Carolina.

Term low birth weight is defined as an infant born at 37 or more completed gestational weeks weighing less than 2500 g and is associated with substantial acute and long-term adverse health effects. Its occurrence has also been linked with exposure to multiple pollutants during different pregnancy time periods across many epidemiological studies (Stieb *et al.*, 2012; Li *et al.*, 2017). A few studies have sought to estimate the critical exposure windows for term birth weight with respect to ambient air pollution exposure but, to our knowledge, none have explored the effect of ignoring potential spatial variability among the lagged health effect parameters on the findings (Darrow *et al.*, 2010; Warren *et al.*, 2013, 2018).

A major consequence of ignoring this spatial variability is that region-specific associations could potentially be masked when grouped with other regions that have different patterns of association. Therefore, we design a simulation study to explore the effect on local risk estimation under different settings for the true variability in the lagged health effect parameters. To offer guidance on determining when consequential spatial variability exists in

a particular application, we investigate the use of two Bayesian model comparison metrics as tools for differentiating spatial variability from spatially constant risk. In combination, this work provides a framework for determining when the variability is sufficiently substantial to warrant use of SpGPCW over existing non-spatial approaches and for estimating spatially varying critical windows when necessary.

2. Data

We work with birth records from 2005–2008 in the North Carolina RDCHCSA (Fig. S1 of the on-line supplementary material). This includes 11 counties in North Carolina (Chatham, Durham, Franklin, Granville, Harnett, Henderson, Johnston, Lee, Orange, Person and Wake) which represent a mix of high population urban centres and less populated rural areas. The inclusion–exclusion criteria for the analysed birth records follow recent work with a similar population (Warren *et al.*, 2019). All singleton live births with a gestational age of at least 37 weeks (determined by the clinical gestational age estimate), without a documented congenital anomaly, with maternal age at delivery between 15 and 50 years inclusively, and occurring in the RDCHCSA between 2005 and 2008, were included in the analysis. This study was approved by the institutional review board at the University of North Carolina at Chapel Hill.

Similarly to Warren *et al.* (2019), we define weekly exposures of PM_{2.5} (24-h average; micrograms per cubic metre) and ozone (8-h maximum; parts per billion (ppb)) for each pregnancy based on calendar dates of pregnancy and residence at delivery (linked with the closest census tract centroid) using the Environmental Protection Agency fused downscaler estimates (Berrocal *et al.* (2010): <https://www.epa.gov/hesc/rsig-related-downloadable-data-files>); all programs that are needed to reproduce the simulation studies in the paper can be found at https://github.com/warrenjl/SpGPCW/tree/master/JRSS-C_Code. Birth records with missing covariate information were excluded from the analyses. Summary information for the study population is displayed in Table S1 of the on-line supplementary material.

3. Methods

3.1. Spatially varying Gaussian process model for critical windows

We model the binary adverse health outcome (e.g. term low birth weight) as a function of individual level covariates or confounders and lagged environmental exposures across pregnancy by using a DLM logistic regression framework. We relax the assumption that has been made in the majority of previous studies that exposure risk is constant across the spatial domain by linking each individual to a subregion of the domain (e.g. county within RDCHCSA) and allowing the lagged risk parameters to vary at that selected spatial level. SpGPCW is given as

$$Y_{ij} | \boldsymbol{\beta}, \boldsymbol{\theta}_i^{\text{ind}} \sim \text{Bernoulli} \{ p_{ij}(\boldsymbol{\beta}, \boldsymbol{\theta}_i) \},$$

$$\ln \left\{ \frac{p_{ij}(\boldsymbol{\beta}, \boldsymbol{\theta}_i)}{1 - p_{ij}(\boldsymbol{\beta}, \boldsymbol{\theta}_i)} \right\} = \mathbf{x}_{ij}^T \boldsymbol{\beta} + \sum_{t=1}^m z_{ij} \{ c_{ij}(t) \} \boldsymbol{\theta}_i(t) \quad (1)$$

where Y_{ij} is equal to 1 if individual j ($j=1, \dots, n_i$) within spatial region i ($i=1, \dots, s$) experiences the outcome and is equal to 0 otherwise, n_i is the total number of individuals residing in spatial region i , s is the total number of unique spatial regions and $\sum_{i=1}^S n_i$ is the total number of individuals in the study; $p_{ij}(\beta, \theta_i)$ is the probability of experiencing the outcome, $\mathbf{x}_{ij}$ is the length p vector of covariates or confounders (including an intercept term) and β describes the association between these variables and $p_{ij}(\beta, \theta_i)$; $z_{ij}(t)$ is the primary exposure of interest averaged over pregnancy period t corresponding to the calendar dates of pregnancy ($c_{ij}(t)$) for individual j in spatial region i (e.g. the average of daily exposures during pregnancy week t for individual j in location i), and $\theta_i(t)$ represents the exposure risk parameter that is specific to spatial region i during pregnancy period t .

Given that we introduce a unique set of m exposure risk parameters at each of the s spatial regions in equation (1), we develop a model for these parameters that allows for data-driven shrinkage across exposure timing and geographic space. Specifically, we use an MCAR model with an exponential time series structure to allow parameters to share information across space and time simultaneously in a computationally convenient way. The exposure risk parameters are modelled as (for $i = 1, \dots, s$)

$$\theta_i | \Theta_{-i}, \eta, \rho, \phi, \sigma_{\theta}^2 \text{ind} \sim \text{MVN} \left\{ \frac{\rho \sum_{j=1}^S w_{ij} \theta_j + (1-\rho) \eta}{\rho \sum_{j=1}^S w_{ij} + 1 - \rho}, \frac{\sigma_{\theta}^2 \Sigma(\phi)}{\rho \sum_{j=1}^S w_{ij} + 1 - \rho} \right\} \quad (2)$$

where θ_i is the vector of m risk parameters that are specific to spatial region i ($\theta_i = \{\theta_i(1), \dots, \theta_i(m)\}^T$), Θ is the collection of all risk parameters ($\Theta = (\theta_1^T, \dots, \theta_s^T)^T$), Θ_{-i} is the Θ -vector with θ_i removed, $\text{MVN}(\mu, \Omega)$ represents the multivariate normal distribution—with mean vector μ and variance–covariance matrix Ω and w_{ij} describes the neighbourhood relationship between spatial regions i and j and is equal to 1 if i and j share a common border and is equal to 0 otherwise. We note that regions are not considered to be neighbours of themselves and we define $w_{ii} = 0$ for all i as is customary.

Model (2) suggests that, conditionally on risk parameter vectors from the other spatial regions, the mean of θ_i is a weighted average of its neighbours' risk parameter vectors and a global mean vector η . The parameter $\rho \in [0, 1]$, whose inclusion results in propriety of the joint distribution of Θ , describes the contribution of spatially structured variability *versus* unstructured variability in the exposure risk parameters. For example, $\rho = 0$ implies a conditional mean of η and covariance of $\sigma_{\theta}^2 \Sigma(\phi)$, suggesting no influence of neighbouring spatial regions. In the other extreme, $\rho = 1$ leads to the improper MCAR model (Besag, 1974; Mardia, 1988) with mean vector $\sum_{j=1}^S w_{ij} \theta_j / \sum_{j=1}^S w_{ij}$ and covariance $\sigma_{\theta}^2 \Sigma(\phi) / \sum_{j=1}^S w_{ij}$, reflecting a strong influence of a region's neighbours.

The strong, positive temporal correlation in risk parameters across pregnancy due to high correlation between exposures is addressed by using $\Sigma(\phi)$ where $\Sigma(\phi)_{t,t'} = \exp\{-\phi|t-t'|\}$ and $\phi > 0$ describes the level of correlation between these parameters. Additionally, the common mean vector η is modelled as

$$\boldsymbol{\eta} = \{\eta(1), \dots, \eta(m)\}^T | \sigma_{\eta}^2, \phi \sim \text{MVN}\{\mathbf{0}_m, \sigma_{\eta}^2 \Sigma(\phi)\} \quad (3)$$

where $\mathbf{0}_m$ is a length m vector with each entry equal to 0. This form allows for temporal smoothness in the risk parameters even in the situation where there is not substantial spatial variability in risk (i.e. $\rho = 0$). The form of model (2) allows for spatial variability in the exposure risk parameter vectors while assuming a separable covariance structure and shared level of temporal correlation through use of the common ϕ -parameter. This permits a computationally tractable MCMC sampling algorithm to be developed.

The conditional form of model (2) can also be expressed jointly such that

$$\boldsymbol{\Theta} | \boldsymbol{\eta}, \rho, \phi, \sigma_{\theta}^2 \sim \text{MVN}(\mathbf{1}_s \otimes \boldsymbol{\eta}, \sigma_{\theta}^2 [\rho(\mathbf{D}_w - \mathbf{W}) + (1 - \rho)\mathbf{I}_s]^{-1} \otimes \Sigma(\phi)) \quad (4)$$

where $\mathbf{1}_s$ is a length s vector with all entries equal to 1, ' $\otimes$ ' represents the Kronecker product, $\mathbf{D}_w$ is a diagonal matrix with (i, i) entry equal to $\sum_{j=1}^S w_{ij}$ $\mathbf{W}$ is the neighbourhood adjacency matrix with (i, j) entry equal to w_{ij} and $\mathbf{I}_s$ is the $s \times s$ identity matrix. This proper version of CAR model was originally developed by Leroux *et al.* (2000) and has been shown to outperform competing spatial random-effect definitions in disease mapping settings (Lee, 2011). Unlike the original form, however, model (4) includes a structured mean vector where each entry in $\boldsymbol{\Theta}$ is centred near the corresponding global gestational week parameter. Additionally, to our knowledge, the multivariate version of the model has not been previously considered within a distributed lag framework.

3.2. Prior distributions

We assign prior distributions to the remaining model parameters and generally favour weakly informative choices given our lack of knowledge regarding the true values and our desire to allow the data to drive the inference rather than our prior beliefs. The regression parameters are specified as $\beta_k \sim \text{IIDN}(0, \sigma_{\beta}^2)$, $k = 1, \dots, p$, where σ_{β}^2 is fixed at a large value. The temporal correlation parameter ϕ is specified as $\phi \sim \text{uniform}(a_{\phi}, b_{\phi})$ where a_{ϕ} and b_{ϕ} are fixed at weakly informative values based on the number of exposure time periods that is considered in the study. The spatial correlation parameter ρ is specified as $\rho \sim \text{uniform}(0, 1)$. Finally, the variance parameters are given independent prior distributions such that $\sigma_{\theta}^2 \sim \text{InverseGamma}(\alpha_{\theta}, \beta_{\theta})$ and $\sigma_{\eta}^2 \sim \text{InverseGamma}(\alpha_{\eta}, \beta_{\eta})$ where α_{θ} , α_{η} , β_{θ} and β_{η} are fixed at weakly informative values.

3.3. Computational efficiency

In appendix A of the on-line supplementary material, we present the model fitting details, including the full conditional distributions for all model parameters. The use of a proper MCAR model offers improvements in computing speed as s becomes large compared with previous distance-based efforts, in addition to being more appropriate given the focus on areal spatial data over point-referenced geostatistical data. Specifically, when updating ρ and ϕ from their respective full conditional distributions, the resulting calculations are less computationally intensive because the inverse of the covariance matrix in model (4) can be

written in closed form with respect to the spatial component (i.e. $\rho(D_w - W) + (1 - \rho)I_s$). This results in the contribution of Θ the full conditional densities having the form

$$\frac{|\rho(D_w - W) + (1 - \rho)I_s|^{m/2}}{\sigma_\theta^{2Sm/2} |\Sigma(\phi)|^{s/2}} \exp \left\{ -\frac{1}{2\sigma_\theta^2} \sum_{k=1}^s \sum_{j=1}^s \kappa_{jk} (\theta_j - \eta)^T \Sigma(\phi)^{-1} (\theta_k - \eta) \right\}$$

where κ_{jk} is the (j, k) th entry of $\rho(D_w - W) + (1 - \rho)I_s$. The only inverse that is involved in this calculation is that of $\Sigma(\phi)$ which is an $m \times m$ matrix where m is generally not large. Using a distance-based Gaussian process approach would require the inverse of an $s \times s$ matrix during each iteration of the algorithm and this would be time consuming as s becomes large. As a result, SpGPCW should be better suited to handle a large number of spatial regions than existing methods are.

4. Simulation study

4.1. Study design

We design a simulation study to investigate the performance of SpGPCW in comparison with the standard non-spatial critical window estimation method under various simulated data settings. We attempt to simulate realistic data sets for analysis by using the RDCHCSA term low birth weight data set as a template, ensuring that our findings are relevant to our data application. Therefore, we consider the same number of spatial regions ($s=11$), the same spatial neighbourhood layout (i.e. the same counties in the RDCHCSA) and the same proportion of individuals residing in each spatial region, when simulating data. To generate a data set with a level of power to detect non-zero risk parameters that is similar to that in our data application while also easing the computational burden of working with a relatively large sample size in a simulation study where repeated model fitting is required across multiple methods, we consider a case-control framework for the simulation analyses. Therefore, for each simulated data set, we use a sample size of four times the number of cases observed in the RDCHCSA term low birth weight application and ensure that approximately 25% of the simulated outcomes result in term low birth weight. This substantially reduces the sample size for each simulated data set while on average providing a similar number of cases to the number that we observed in the data application.

We simulate data from equation (1) where $\mathbf{x}_{ij} = 1$ for all i and j such that only an intercept parameter is included. We adjust the intercept for each simulation setting to ensure that the approximate correct number of cases is obtained (i.e. 25%). In terms of the primary exposures, we randomly select actual exposure trajectories (i.e. a set of exposures across the entire pregnancy for a randomly selected woman) from our larger RDCHCSA analysis data set, without replacement, and assign them to each simulated individual to ensure that we are working with realistic exposures with respect to distribution and correlation across pregnancy. A unique set of exposures is sampled from the larger pool of women for each simulated data set. As in the data application (Section 5.1), we scale the exposures separately at each week of gestation by dividing by the interquartile range IQR.

We consider three simulation settings that involve different assumptions for the risk parameters:

- a. setting S1 (single set of risk parameters), where $\theta_i \equiv \eta$ for all i where η is generated from equation (3);
- b. setting S2 (spatially varying risk parameters, high spatial correlation), where θ_i follows the models in equations (2) and (3) with ρ fixed at a realistic value;
- c. setting S3 (spatially varying risk parameters, no spatial correlation), where θ_i follows the models in equations (2) and (3) with ρ fixed at 0.

Across all settings, we set those simulated risk parameters which are less than 0.04 in absolute value (based on results seen in the data application of Section 5) to 0 to resemble more closely the data-generating process of Warren *et al.* (2019). This results in critical window sets where potentially multiple clusters of weekly periods that vary in magnitude, sign, timing and duration are thought to have an effect on pregnancy outcomes whereas others do not (Wang *et al.*, 2018). Examples of simulated risk parameters from each simulation setting are shown in Fig. S2 of the on-line supplementary material.

Setting S1 assumes that there is a single set of risk parameters that is shared by every individual, regardless of spatial location. Setting S2 assumes that there is spatial variability in the risk parameters and that the degree of spatial correlation between them is very high, resulting in less pronounced differences in risk across the spatial regions. Setting S3 assumes spatial variability in the risk parameters but specifies no spatial correlation, resulting in more distinct differences in risk across space.

In all settings, we simulate a new set of risk parameters for every simulated data set and centre them at zero in each spatial region to resemble more closely the risk parameters that are estimated in the data application. For simulating risk parameters from equations (2) and (3), the hyperparameters of the models are fixed at values estimated by SpGPCW applied to the term low birth weight PM_{2.5} data set. We use the posterior means of σ_θ^2 , σ_η^2 , ρ and ϕ for simulating data. See Table S2 of the on-line supplementary material for the specific values for these parameters. To test the ability of SpGPCW to estimate ϕ accurately, we repeat the data simulation process from setting S3 using $\phi=0.29$ and $\phi=0.69$ (100 data sets for each value). We then apply SpGPCW to these data sets and monitor the estimation of ϕ . For the primary study, we simulate 100 data sets from each setting for a total of 300 for analysis.

We look to compare the performance of SpGPCW with that of the standard method of analysis under each of the simulation settings proposed. We define the standard method as a Gaussian process for critical window identification, ‘GPCW’. Versions of this method have been used in various past studies to explore exposure risk across pregnancy under the assumption that risk is constant across space (Warren *et al.*, 2012b, 2013, 2017). GPCW is given as

$$Y_{ij} | \beta, \eta \stackrel{\text{ind}}{\sim} \text{Bernoulli}\{p_{ij}(\beta, \eta)\}$$

with

$$\ln \left\{ \frac{p_{ij}(\boldsymbol{\beta}, \boldsymbol{\eta})}{1 - p_{ij}(\boldsymbol{\beta}, \boldsymbol{\eta})} \right\} = \mathbf{x}_{ij}^T \boldsymbol{\beta} + \sum_{t=1}^m z_{ij} \{c_{ij}(t)\} \eta(t)$$

where the model for $\boldsymbol{\eta}$ is given in equation (3). Therefore, GPCW matches setting S1. As a result, we hypothesize that GPCW will perform well in setting S1 where there is a single set of risk parameters but will increasingly struggle in settings S2 and S3 as the spatial variability in risk increases. SpGPCW in contrast should be able to perform well in each setting since it can technically collapse to GPCW when $\sigma_{\theta}^2 \approx 0$ and $\rho \approx 0$.

Each competing method is fitted to each simulated data set for a total of 25000 posterior samples (after removing the first 5000 as a burn-in period) where we further thin the samples by a factor of 5 to reduce auto-correlation in the final set of posterior samples. The prior distributions are described in Section 3.2 with specific settings detailed in Section 5.1. Data are then collected from each fit so that we can formally compare the performances of the methods. We first look to determine which method produces the best inference for the risk parameters in each simulation setting. Therefore, we estimate the average mean-squared error MSE of the posterior mean estimators and the average empirical coverage EC of the 95% quantile-based, equal-tailed credible intervals (CIs) (averaged across all risk parameters). The average MSE for a particular simulation setting and method is estimated as

$$\frac{1}{100sm} \sum_{i=1}^{100} \left[\sum_{j=1}^s \sum_{t=1}^m \{ \hat{\theta}_{ij}(t) - \theta_{ij}(t) \}^2 \right]$$

where $\hat{\theta}_{ij}(t)$ is the posterior mean of the risk parameter for simulated data set i , at spatial region j , during pregnancy time period t and $\theta_{ij}(t)$ is the corresponding true parameter. For GPCW, $\hat{\theta}_{ij}(t) = \hat{\eta}_i(t)$ for all j since the model assumes the same risk across space. Similarly, in setting S1, $\theta_{ij}(t) = \eta_i(t)$. The average EC for a particular simulation setting and method is estimated as

$$\frac{1}{100sm} \sum_{i=1}^{100} \left[\sum_{j=1}^s \sum_{t=1}^m I \{ \theta_{ij}(t) \in \text{CI}_{ij}(t) \} \right]$$

where $\text{CI}_{ij}(t)$ is the CI for simulated data set i , spatial region j and pregnancy period t .

Additionally, we quantify the ability of each method to identify the true critical windows correctly by estimating average sensitivity and specificity. For a single analysed data set, we compare the true critical window set (i.e. those not equal to zero) with those estimated after fitting the model (i.e. those whose 95% CIs exclude zero) to estimate sensitivity and specificity in each spatial region, and we average the respective values across all spatial regions. We then average these estimates over the 100 simulated data sets. Finally, we are also interested in determining whether the Watanabe–Akaike information criterion WAIC (Watanabe, 2010) and/or a posterior predictive model comparison metric PPMC (Laud and

Ibrahim, 1995) can be used to identify data sets where there is true spatial variability in the exposure risk parameters. Therefore, for each fitted model, we calculate the proportion of times that WAIC and PPMC correctly identify the data-generating setting. For computing PPMC, we use the binomial deviance discrepancy function with continuity correction as detailed in Gelfand and Ghosh (1998).

4.2. Study results

The results from the simulation study are displayed in Table 1. Overall, SpGPCW performs well in each simulated data setting as expected. In terms of average MSE and sensitivity, SpGPCW is never meaningfully outperformed by GPCW, even in setting S1 where GPCW is technically the correct data-generating model. GPCW struggles in terms of MSE, EC and specificity as the variability in the exposure risk parameters increases, whereas SpGPCW consistently offers improvements. Results suggest that WAIC can be used to detect true variability in the risk parameters where it selected the correct model 98%, 91% and 100% of the times in settings S1, S2 and S3 respectively. PPMC appears to favour the more complex model in most settings and as a result may not be as useful in this context. Overall, SpGPCW is robust in offering improved risk parameter estimation and in identifying true critical window sets, and it generally outperforms GPCW under any level of variability in the risk parameters.

In Fig. S3 of the on-line supplementary material, we display histograms of posterior means for the covariance hyperparameters based on fitting SpGPCW to data simulated from setting S3 under the three settings for the true value of ϕ (0:0002, 0:29, 0:69). Results suggest that SpGPCW is also performing well with respect to estimating these parameters.

5. Term low birth weight in the Raleigh–Durham–Chapel Hill combined statistical area, 2005–2008

5.1. Study design

We use SpGPCW to analyse the risk of term low birth weight with respect to weekly PM_{2.5} and ozone exposure during pregnancy, and to determine whether this risk differs geographically over the 11-county RDCHCSA in 2005–2008. To investigate the potential of spatial variability in risk more formally, we also fit GPCW and compare the models by using WAIC and PPMC, where WAIC was shown to be a useful tool in identifying true variability in Section 4.2 (see Table 1). The methods are fitted separately for each pollutant.

With respect to covariates or confounders in $\mathbf{x}_{ij}$, we include indicators for maternal smoking during pregnancy and diabetes status (which includes both gestational and pregestational), parity, maternal age category ([15,25), [25,30), [30,35), [35,50] years), maternal education level (less than high school, high school, greater than high school), maternal race and ethnicity (white non-Hispanic, black non-Hispanic, Hispanic), sex of the child, marital status, trimester that prenatal care began (first *versus* second or third), the latitude, longitude and their interaction of the residence at delivery, and monthly and yearly time indicator variables to account for seasonal and long-term time trends respectively. Continuous covariates were standardized by subtracting the mean and dividing by the standard deviation

before model fitting. From equation (1), we set $s = 11$ to represent the 11 counties in the study and $m = 37$ since we consider only the set of exposures that was included for every woman in the study. From Section 3.2, we specify $\alpha_\theta = \alpha_\eta = 3$, $\beta_\theta = \beta_\eta = 2$, $\sigma_\beta = 100$, $a_\phi = -\ln.(0.9999)/36$ and $b_\phi = -\ln.(0.0001)$. The values for a_ϕ and b_ϕ were selected on the basis of the number of lagged periods that was included in the study ($m = 37$) and the desire to allow for the possibility of strong spatial correlation (0.9999) at the largest lag distance (36) and weak spatial correlation (0.0001) at the smallest lag difference (1) *a priori*. These specifications were set and we then solved for a_ϕ and b_ϕ . In a sensitivity analysis, we also consider $\phi \sim \text{Gamma}(.0:10, 0:10)$ as a weakly informative alternative prior choice.

All models were fitted for 250000 iterations with the first 150000 removed as a burn-in period. We thinned the remaining samples by a factor of 10 to reduce auto-correlation in the posterior samples, but also to reduce the dimensionality of the results that need to be saved and further processed. Convergence was assessed through visual inspection of trace plots and calculation of the Geweke convergence diagnostic (Geweke, 1991) for each individual parameter. Across all models, none of the individual parameters failed this diagnostic test after Bonferroni adjustment for multiple comparisons. We also calculated the effective sample size for each individual parameter to ensure that we had collected an appropriate number of posterior samples post convergence to make accurate posterior inference. Table S3 of the on-line supplementary material shows specific convergence details for all models.

Exposures were scaled separately at each week of gestation by dividing the original exposure by the corresponding IQR. This results in an IQR-interpretation for the risk parameters on the odds ratio (OR) scale (for example, an IQR increase in exposure at time period t results in an $\exp\{\theta_i(t)\}$ OR increase in risk for spatial region i). The weekly IQRs ranged from 5.66 to 5.87 $\mu\text{g m}^{-3}$ and from 22.58 to 23.25 ppb for $\text{PM}_{2.5}$ and ozone respectively.

5.2. Study results

Table 2 displays the posterior inference on the OR scale for the covariates or confounders that were included in the study from the $\text{PM}_{2.5}$ SpGPCW data application. The corresponding results from the ozone model fit are nearly identical and are not displayed as a result. In Table S2 of the on-line supplementary material, we present posterior inference from SpGPCW for the covariance parameters and hyperparameters for each pollutant. The results are not overly sensitive to the prior choice for ϕ as seen by the similar posterior summaries for ϕ across both prior distributions.

From Table 2 it is clear that smoking during pregnancy, being a non-Hispanic black or other race-ethnicity mother (compared with non-Hispanic white), having a female child and being unmarried are potential risk factors for term low birth weight, whereas increased parity and having more than a high school education appear to be protective.

Based on the simulation study findings, the WAIC- (and PPMC-) results from the fitted models (see Table S4 of the on-line supplementary material) suggest that there is substantial spatial variability in the lagged risk parameters for each pollutant. Table S2 of the on-line supplementary material indicates that on the basis of the posterior distribution of ρ for each

pollutant this excess variability may be spatially structured (posterior mean 0.77 for PM_{2.5} and 0.76 for ozone), though this finding should be interpreted with caution given the small number of spatial regions in the analysis and the difficulty in learning about ρ as a result.

In Fig. 1, we present results from GPCW and SpGPCW for the PM_{2.5} data application. Posterior means (OR scale) along with 95% CIs are displayed. We include only those counties from the RDCHCSA that had at least one 95% CI that excluded 1 (in either direction) with the remaining counties displayed in Fig. S4 of the on-line supplementary material. The GPCW results suggest that there is a protective effect of increased exposure to PM_{2.5} during the first month of pregnancy. However, when SpGPCW is used to investigate risk at the county level, the significance of this effect is reduced. In addition, signal that has been masked through the use of GPCW is revealed when the individual county level risk is analysed. In total, 13 pregnancy weeks with 95% CIs that exclude 1 are missed when GPCW is used to estimate critical windows. This increased risk is present in the late first and early second trimesters across the various counties. In

Fig. 2, the ORs are mapped across the RDCHCSA and exposure time for the relevant pregnancy periods. For the ozone results in Fig. S5 of the on-line supplementary material, GPCW suggests no increased risk of term low birth weight though SpGPCW identifies elevated risk in Johnston county during pregnancy week 8 (OR 1.04; 95% CI 1.00–1.09). The remaining county results from the ozone analysis are displayed in Fig. S6 of the on-line supplementary material.

6. Discussion

Through simulation and data application, we have shown the importance of accounting for spatially varying lagged health effect parameters for accurately estimating and quantifying uncertainty in risk at more localized spatial scales. Allowing for spatially varying lagged parameters will be increasingly important in settings where the sample size is small in some or all of the discrete spatial regions. SpGPCW benefits from being able to collapse to a non-spatial version of the model when $\rho=0$ and even further to a single set of risk parameters shared across all spatial regions (GPCW) when $\rho=0$ and $\sigma_{\theta}^2=0$. WAIC was shown to be a useful tool with respect to determining whether spatial variability is present in the risk parameters. In our application of SpGPCW, it is difficult to differentiate spatial and non-spatial variability in the risk parameters because of the relatively small number of spatial regions that were considered ($s=11$) and the large number of individuals in each spatial region. Although the posterior distribution of ρ is slightly shifted upwards in both applications (Table S2 of the on-line supplementary material), suggesting spatial correlation in the risk parameters, we expect to see similar risk parameter results when a model that ignores spatial correlation is used (e.g. $\rho=0$). However, as the number of spatial regions increases and/or the number of individuals within each spatial region decreases, we expect ρ to play a more important role in estimation of the risk parameters as it controls the amount of spatial shrinkage across regions.

Additionally, as the number of spatial regions increases, SpGPCW will be easier to work with computationally than a distance-based Gaussian process approach because of the use of

an MCAR model in the framework. In our application of the model, the computational difference between these methods would probably be negligible because of the relatively small number of spatial regions being considered ($s=11$). With respect to computation, the difficulty in this application arises from the large number of people who were included in the study. Future work with these models should closely investigate the computational gains of SpGPCW and how this varies on the basis of the number of unique spatial regions and the individual sample sizes within each region.

When used to estimate the association between air pollution exposure and term low birth weight in the RDCHCSA, GPCW did not identify any positive statistically significant associations because it lacks the flexibility to model spatial region-specific lagged health effect parameters. When using SpGPCW, a somewhat consistent picture emerged of elevated risk to PM_{2.5}-exposure during late first trimester through early second trimester (weeks 9–17). This finding is in agreement with a recent critical window study of the effect of PM₁₀ on term low birth weight in Connecticut (Warren *et al.*, 2017) which found elevated associations in weeks 16–18, though it differs when compared with a Texas-based study on PM_{2.5}-exposure that found increased associations during weeks 20–23 (Warren *et al.*, 2013).

As with much of the previous work in this area, a limitation is the inability to account for residential mobility during pregnancy. Here, we use residential address at delivery when assigning exposures as well as the use of ambient concentrations to represent personal exposure. Previous work has shown that pregnant women who change residences in early pregnancy tend to move only short distances. Critical window estimation has been shown to be robust to this type of exposure misclassification (Bell and Belanger, 2012; Warren *et al.*, 2017) and we do not expect that it would substantially bias our results. Our use of modelled estimates of air pollution concentrations from the Environmental Protection Agency's community multiscale air quality downscaler model represents a strength with respect to exposure assignment as it covers spatial and temporal areas where air pollution monitors are not present and includes spatiotemporally varying bias correction for measured air pollution concentrations.

In summary, we apply SpGPCW and GPCW to investigate how ignoring spatial variability in lagged health effect parameters can affect the estimation of critical exposure windows in studies of air pollution and term low birth weight. Spatial variation in risk may be due to exposure characteristics (e.g. spatially varying PM_{2.5}-composition, spatial differences in exposure measurement error) and/or residual confounding (e.g. effects of not controlling for other pollutants that are more spatially heterogeneous or inclusion of sociodemographic variables that might not be accurate proxies for the contextual confounders that they are meant to represent). The demonstrated ability of SpGPCW to account for spatially varying lagged parameters will enable more accurate risk estimation in epidemiologic studies, particularly in the presence of these modelling issues. Future work to determine carefully the size of the study domain before using methods to identify or estimate critical windows along with use of SpGPCW to understand spatial variability in the estimated risk better will lead to an improved understanding of exposure timing and its effect on health. Future methods work could also focus on extending the methodology to the variable-selection setting where spatial correlation will need to be considered and on using alternative time series structures

for modelling the distributed lags across time (e.g. auto-regressive models). An R package (SpGPCW) has been developed to implement the new method with code available from <https://github.com/warrenjl/SpGPCW> (including code to reproduce the simulation study). This offers an opportunity for collaboration with researchers who are interested in extending the method and/or improving computational aspects of the algorithm.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgements

This work was supported by the National Institutes of Environmental Health Sciences (grant R01 NIEHS ES028346). The authors thank the North Carolina State Center for Health Statistics for providing the birth data. The views expressed do not necessarily represent the views or policies of the US Environmental Protection Agency.

References

- Baek J, Sanchez-Vaznaugh EV and Sanchez BN (2016) Hierarchical distributed-lag models: exploring varying geographic scale and magnitude in associations between the built environment and health. *Am. J. Epidemiol.* 183, 583–592.
- Banerjee S, Carlin BP and Gelfand AE (2014) *Hierarchical Modeling and Analysis for Spatial Data*. Boca Raton: CRC Press.
- Bell ML and Belanger K (2012) Review of research on residential mobility during pregnancy: consequences for assessment of prenatal environmental exposures. *J. Expos. Sci. Environ. Epidemiol.* 22, 429–438.
- Bello GA, Arora M, Austin C, Horton MK, Wright RO and Gennings C (2017) Extending the distributed lag model framework to handle chemical mixtures. *Environ. Res.* 156, 253–264. [PubMed: 28371754]
- Berrocal VJ, Gelfand AE and Holland DM (2010) A spatio-temporal downscaler for output from numerical models. *J. Agric. Biol. Environ. Statist.* 15, 176–197.
- Besag J (1974) Spatial interaction and the statistical analysis of lattice systems. *J. R. Statist. Soc. B.* 36, 192–225.
- Bobb JF, Valeri L, Claus Henn B, Christiani DC, Wright RO, Mazumdar M, Godleski JJ and Coull BA (2014) Bayesian kernel machine regression for estimating the health effects of multi-pollutant mixtures. *Biostatistics*, 16, 493–508. [PubMed: 25532525]
- Chang HH, Warren JL, Darrow LA, Reich BJ and Waller LA (2015) Assessment of critical exposure and outcome windows in time-to-event analysis with application to air pollution and preterm birth study. *Biostatistics*, 16, 509–521. [PubMed: 25572998]
- Chen Y-H, Mukherjee B and Berrocal VJ (2019) Distributed lag interaction models with two pollutants. *Appl. Statist.* 68, 79–97.
- Darrow LA, Klein M, Strickland MJ, Mulholland JA and Tolbert PE (2010) Ambient air pollution and birth weight in full-term infants in Atlanta, 1994–2004. *Environ. Health Perspect.* 119, 731–737.
- Dominici F, Peng RD, Bell ML, Pham L, McDermott A, Zeger SL and Samet JM (2006) Fine particulate air pollution and hospital admission for cardiovascular and respiratory diseases. *J. Am. Med. Ass.* 295, 1127–1134.
- Gasparrini A, Armstrong B and Kenward MG (2010) Distributed lag non-linear models. *Statist. Med.* 29, 2224–2234.
- Gelfand AE and Ghosh SK (1998) Model choice: a minimum posterior predictive loss approach. *Biometrika*, 85, 1–11.
- Geweke J (1991) *Evaluating the Accuracy of Sampling-based Approaches to the Calculation of Posterior Moments*. Minneapolis: Federal Reserve Bank of Minneapolis.

- Heaton MJ, Datta A, Finley AO, Furrer R, Guinness J, Guhaniyogi R, Gerber F, Gramacy RB, Hammerling D, Katzfuss M, Lindgren F, Nychka DW, Sun F and Zammit-Mangion A (2018) A case study competition among methods for analyzing large spatial data. *J. Agric. Biol. Environ. Statist*, 24, 398–425.
- Heaton MJ and Peng RD (2013) Extending distributed lag models to higher degrees. *Biostatistics*, 15, 398–412. [PubMed: 23990524]
- Hu H, Téllez-Rojo MM, Bellinger D, Smith D, Ettinger AS, Lamadrid-Figueroa H, Schwartz J, Schnaas L, Mercado-García A and Hernandez-Avila M (2006) Fetal lead exposure at each stage of pregnancy as a predictor of infant mental development. *Environ. Hlth Perspect*, 114, 1730–1735.
- Kavanagh L, Lee D and Pryce G (2016) Is poverty decentralizing?: Quantifying uncertainty in the decentralization of urban poverty. *Ann. Am. Ass. Geog.*, 106, 1286–1298.
- Laud PW and Ibrahim JG (1995) Predictive model selection. *J. R. Statist. Soc. B*, 57, 247–262.
- Lee D (2011) A comparison of conditional autoregressive models used in Bayesian disease mapping. *Spatl Spatio-temp. Epidem*, 2, 79–89.
- Leon Hsu H-H, Mathilda Chiu Y-H, Coull BA, Kloog I, Schwartz J, Lee A, Wright RO and Wright RJ (2015) Prenatal particulate air pollution and asthma onset in urban children: identifying sensitive windows and sex differences. *Am. J. Resp. Crit. Care Med*, 192, 1052–1059. [PubMed: 26176842]
- Leroux BG, Lei X and Breslow N (2000) Estimation of disease rates in small areas: a new mixed model for spatial dependence In *Statistical Models in Epidemiology, the Environment, and Clinical Trials* (eds Halloran EM and Berry D), pp. 179–191. New York: Springer.
- Li X, Huang S, Jiao A, Yang X, Yun J, Wang Y, Xue X, Chu Y, Liu F, Liu Y, Ren M, Chen X, Li N, Lu Y, Mao Z, Tian L and Xiang H (2017) Association between ambient fine particulate matter and preterm birth or term low birth weight: an updated systematic review and meta-analysis. *Environ. Polln*, 227, 596–605.
- Liu SH, Bobb JF, Lee KH, Gennings C, Claus Henn B, Bellinger D, Austin C, Schnaas L, Tellez-Rojo MM, Hu H, Wright RO, Arora M and Coull BA (2017) Lagged kernel machine regression for identifying time windows of susceptibility to exposures of complex mixtures. *Biostatistics*, 19, 325–341.
- Mardia K (1988) Multi-dimensional multivariate Gaussian Markov random fields with application to image processing. *J. Multiv. Anal*, 24, 265–284.
- National Institute of Environmental Health Sciences (2012) 2012–2017 strategic plan—advancing science, improving health; a plan for environmental health research Publication NIH 12–7935. National Institute of Environmental Health Sciences, Bethesda.
- Peel JL, Tolbert PE, Klein M, Metzger KB, Flanders WD, Todd K, Mulholland JA, Ryan PB and Frumkin H (2005) Ambient air pollution and respiratory emergency department visits. *Epidemiology*, 16, 164–174. [PubMed: 15703530]
- Samet JM, Zeger SL, Dominici F, Currier F, Coursac I, Dockery DW, Schwartz J and Zanobetti A (2000) The national morbidity, mortality, and air pollution study: Part II, morbidity and mortality from air pollution in the United States. *Res. Rep. Hlth Eff. Inst*, 94, 5–79.
- Sanchez BN, Hu H, Litman HJ and Téllez-Rojo MM (2010) Statistical methods to study timing of vulnerability with sparsely sampled data on environmental toxicants. *Environ. Hlth Perspect*, 119, 409–415.
- Schwartz J (2000) The distributed lag between air pollution and daily deaths. *Epidemiology*, 11, 320–326. [PubMed: 10784251]
- Stieb DM, Chen L, Eshoul M and Judek S (2012) Ambient air pollution, birth weight and preterm birth: a systematic review and meta-analysis. *Environ. Res*, 117, 100–111. [PubMed: 22726801]
- Wang Q, Benmarhnia T, Zhang H, Knibbs LD, Sheridan P, Li C, Bao J, Ren M, Wang S, He Y, Zhang Y, Zhao Q and Huang C (2018) Identifying windows of susceptibility for maternal exposure to ambient air pollution and preterm birth. *Environ. Int*, 121, 317–324. [PubMed: 30241019]
- Warren J, Fuentes M, Herring A and Langlois P (2012a) Spatial-temporal modeling of the association between air pollution exposure and preterm birth: identifying critical windows of exposure. *Biometrics*, 68, 1157–1167. [PubMed: 22568640]

- Warren J, Fuentes M, Herring A and Langlois P (2012b) Bayesian spatial-temporal model for cardiac congenital anomalies and ambient air pollution risk assessment. *Environmetrics*, 23, 673–684. [PubMed: 23482298]
- Warren JL, Fuentes M, Herring AH and Langlois PH (2013) Air pollution metric analysis while determining susceptible periods of pregnancy for low birth weight. *Int. Schol. Res. Nots Obstet. Gyn*, article 387452.
- Warren JL, Kong W, Luben TJ and Chang HH (2019) Critical window variable selection: estimating the impact of air pollution on very preterm birth. *Biostatistics* to be published, doi 10.1093/biostatistics/kxz006.
- Warren JL, Son J-Y, Pereira G, Leaderer BP and Bell ML (2018) Investigating the impact of maternal residential mobility on identifying critical windows of susceptibility to ambient air pollution during pregnancy. *Am. J. Epidem*, 187, 992–1000.
- Watanabe S (2010) Asymptotic equivalence of Bayes cross validation and widely applicable information criterion in singular learning theory. *J. Mach. Learn. Res*, 11, 3571–3594.
- Welty LJ and Zeger SL (2005) Are the acute effects of particulate matter on mortality in the National Morbidity, Mortality, and Air Pollution Study the result of inadequate control for weather and season?: A sensitivity analysis using flexible distributed lag models. *Am. J. Epidem*, 162, 80–88.
- Wilson A, Chiu Y-HM, Hsu H-HL, Wright RO, Wright RJ and Coull BA (2017) Bayesian distributed lag interaction models to identify perinatal windows of vulnerability in children's health. *Biostatistics*, 18, 537–552. [PubMed: 28334179]
- Wilson A, Chiu Y-HM, Hsu H-HL, Wright RO, Wright RJ and Coull BA (2019) Kernel machine and distributed lag models for assessing windows of susceptibility to mixtures of time-varying environmental exposures in children's health studies Preprint arXiv:1904.12417. Department of Statistics, Colorado State University, Fort Collins.
- Zanobetti A and Schwartz J (2009) The effect of fine and coarse particulate air pollution on mortality: a national analysis. *Environ. Hlth Perspect*, 117, 898–903.
- Zanobetti A, Wand M, Schwartz J and Ryan L (2000) Generalized additive distributed lag models: quantifying mortality displacement. *Biostatistics*, 1, 279–292. [PubMed: 12933509]

Fig. 1. Posterior mean and 95% CI results from the term low birth weight and $PM_{2.5}$ (24-h average) exposure analysis in the RDCHCSA, 2005–2008 (results are based on an interquartile range increase in weekly exposure) (weeks identified as part of the critical window set (harmful)): (a) GPCW; (b) Durham; (c) Granville; (d) Johnston; (e) Vance

Fig. 2. Posterior mean OR results from the term low birth weight and PM_{2.5} (24-h average) exposure analysis in the RDCHCSA, 2005–2008 (results based on an interquartile range increase in weekly exposure):(a) week 9; (b) week 10; (c) week 11; (d) week 12; (e) week 13; (f) week 14; (g) week 15; (h) week 16; (i) week 17

Table 1.

Average MSE, average empirical coverage, average sensitivity, average specificity, WAIC and PPMC simulation study results across all simulation settings and methods[‡]

Metric	Method	Results for setting S1 (non-spatial)	Results for setting S2 (spatial, high correlation)	Results for setting S3 (spatial, no correlation)
MSE [‡]	GPCW	3.06 (0.14)	5.83 (0.21)	9.03 (0.26)
	SpGPCW	3.20 (0.14)	4.98 (0.15)	6.34 (0.15)
Coverage	GPCW	0.95 (0.00)	0.88 (0.01)	0.78 (0.01)
	SpGPCW	0.94 (0.00)	0.92 (0.01)	0.94 (0.00)
Sensitivity	GPCW	0.50 (0.03)	0.34 (0.03)	0.41 (0.04)
	SpGPCW	0.44 (0.03)	0.43 (0.03)	0.34 (0.02)
Specificity	GPCW	0.98 (0.00)	0.67 (0.03)	0.61 (0.04)
	SpGPCW	0.97 (0.00)	0.96 (0.00)	0.97 (0.00)
WAIC	Proportion correct	0.98 (0.01)	0.91 (0.03)	1.00 (0.00)
PPMC	Proportion correct	0.03 (0.02)	1.00 (0.00)	1.00 (0.00)

[‡]Sensitivity and specificity correspond to the identification of true critical window sets. Standard errors are presented in parentheses.

[‡]Multiplied by 10000 for presentation purposes.

Table 2.

Maternal and child covariate results from the term low birth weight and PM_{2.5} (24-h average) exposure analysis in the RDCHCSA, 2005–2008[†]

Parameter	Mean	Standard deviation	Quantiles	
			0.025	0.975
Smoked during pregnancy				
Yes <i>versus</i> no	2.30	0.17	1.98	2.66
Diabetes status				
Yes <i>versus</i> no	1.12	0.21	0.75	1.56
Parity				
> 0 <i>versus</i> 0	0.59	0.03	0.54	0.66
Age (years)				
[25, 30] <i>versus</i> [15,25)	0.94	0.06	0.83	1.07
[30, 35] <i>versus</i> [15,25)	0.95	0.07	0.81	1.10
[35,50] <i>versus</i> [15,25)	1.09	0.10	0.91	1.29
Education				
High school <i>versus</i> less than high school	0.94	0.06	0.82	1.07
More than high school <i>versus</i> less than high school	0.59	0.05	0.51	0.69
Race-ethnicity				
Black, non-Hispanic <i>versus</i> white, non-Hispanic	2.16	0.14	1.91	2.45
Hispanic <i>versus</i> white, non-Hispanic	0.93	0.08	0.79	1.09
Other <i>versus</i> white, non-Hispanic	1.87	0.21	1.49	2.29
Sex of child				
Female <i>versus</i> male	1.73	0.08	1.57	1.91
Marital status				
Unmarried <i>versus</i> married	1.16	0.07	1.02	1.30
Trimester when prenatal care began				
Second or third <i>versus</i> first	1.04	0.07	0.91	1.17

[†]Model also adjusted for temporal trends, seasonality and spatial location. The posterior inference on the OR scale is presented.