

# **Lagged kernel machine regression for identifying time windows of susceptibility to exposures of complex mixtures**

SHELLEY H. LIU\*

*Department of Population Health Science and Policy, Icahn School of Medicine at Mount Sinai, 1425  
Madison Avenue, New York, NY 10029, USA*

shelley.liu@mountsinai.org

JENNIFER F. BOBB

*Biostatistics Unit, Kaiser Permanente Washington Health Research Institute, 1730 Minor Avenue,  
Seattle, WA 98101, USA*

KYU HA LEE

*Epidemiology and Biostatistics Core, The Forsyth Institute, 245 First Street, Cambridge, MA 02142, USA*

CHRIS GENNINGS

*Department of Environmental Medicine and Public Health, Icahn School of Medicine at Mount Sinai, 17  
E. 102nd Street, New York, NY 10029, USA*

BIRGIT CLAUS HENN

*Department of Environmental Health, Boston University School of Public Health, 715 Albany Street,  
Boston, MA 02118, USA*

DAVID BELLINGER

*Department of Environmental Health, Harvard T.H. Chan School of Public Health, 665 Huntington  
Avenue, Boston, MA 02115, USA*

CHRISTINE AUSTIN

*Department of Environmental Medicine and Public Health, Icahn School of Medicine at Mount Sinai, 17  
East 102nd Street, New York, NY 10029, USA*

LOURDES SCHNAAS

*Division of Research in Community Interventions, National Institute of Perinatology, Montes Urales  
800, Lomas Virreyes, CP 11000, CDMX, México*

\*To whom correspondence should be addressed.

MARTHA M. TELLEZ-ROJO

*Center for Nutrition and Health Research, National Institute of Public Health,  
Universidad No. 655 Colonia Santa María, Ahuacatitlán, Cerrada Los Pinos y Caminera,  
C.P. 62100, Cuernavaca, Morelos, México*

HOWARD HU

*Dalla Lana School of Public Health, University of Toronto, 155 College Street,  
Toronto, ON, M5T 3M7, Canada*

ROBERT O. WRIGHT, MANISH ARORA

*Department of Environmental Medicine and Public Health, Icahn School of Medicine at Mount Sinai,  
17 E. 102nd Street, New York, NY 10029, USA*

BRENT A. COULL

*Department of Biostatistics, Harvard T.H. Chan School of Public Health, 665 Huntington Avenue,  
Boston, MA 02115, USA*

#### SUMMARY

The impact of neurotoxic chemical mixtures on children's health is a critical public health concern. It is well known that during early life, toxic exposures may impact cognitive function during critical time intervals of increased vulnerability, known as windows of susceptibility. Knowledge on time windows of susceptibility can help inform treatment and prevention strategies, as chemical mixtures may affect a developmental process that is operating at a specific life phase. There are several statistical challenges in estimating the health effects of time-varying exposures to multi-pollutant mixtures, such as: multi-collinearity among the exposures both within time points and across time points, and complex exposure–response relationships. To address these concerns, we develop a flexible statistical method, called lagged kernel machine regression (LKMR). LKMR identifies critical exposure windows of chemical mixtures, and accounts for complex non-linear and non-additive effects of the mixture at any given exposure window. Specifically, LKMR estimates how the effects of a mixture of exposures change with the exposure time window using a Bayesian formulation of a grouped, fused lasso penalty within a kernel machine regression (KMR) framework. A simulation study demonstrates the performance of LKMR under realistic exposure-response scenarios, and demonstrates large gains over approaches that consider each time window separately, particularly when serial correlation among the time-varying exposures is high. Furthermore, LKMR demonstrates gains over another approach that inputs all time-specific chemical concentrations together into a single KMR. We apply LKMR to estimate associations between neurodevelopment and metal mixtures in Early Life Exposures in Mexico and Neurotoxicology, a prospective cohort study of child health in Mexico City.

*Keywords:* Bayesian analysis; Environmental epidemiology; Hierarchical models; Statistical methods in epidemiology.

#### 1. INTRODUCTION

A critical public health concern is the impact of neurotoxic chemicals on children's cognitive development. There is a large body of literature on the impact of exposure to individual chemicals, such as lead, on neurodevelopment (Bellinger, 2008; Hu and others, 2006; Sanchez and others, 2011). However, exposure to chemical mixtures, rather than to individual chemicals, are more reflective of real-world scenarios.

Accordingly, the National Institute of Environmental Health Sciences (NIEHS) has placed a priority on quantification of the health impacts of exposure to environmental mixtures (Carlin *and others*, 2013).

Methods to address the analysis of environmental mixtures are needed if we are to understand the role that chemicals play in child development. However, the analysis of combinations and doses of chemicals that are toxic to children requires methods that can address the role of biological development, in which rapid changes in gene expression, tissue growth, and cell differentiation alter the susceptibility of a child's body to toxic chemicals.

Estimation of the health effects of metal mixtures on neurodevelopment can be complex. In particular, the shape of the exposure–response relationship may exhibit both non-linearity and non-additivity. The effect of some metals, such as trace elements like manganese, can be non-linear as they are essential nutrients at low doses but neurotoxic at high exposure levels. These dual roles can result in an inverted-*u* relationship with neurodevelopment (Claus Henn *and others*, 2010). Moreover, existing work on metal mixtures provides evidence of interactions between individual metals. For instance, summarizing the expanding body of literature, Claus Henn *and others* (2014) reported increased lead toxicity in the presence of higher levels of manganese, arsenic, mercury, and cadmium. In theory, interaction may occur at specific dose ranges and may not occur at other doses, making the need for statistical methods that can capture all these properties critical.

Another layer of complexity in the identification of environmental effects on children's health is that health effects can be highly dependent on exposure timing. There exist many sequential developmental processes in early life, as development is unidirectional and specifically-timed, occurring as a complex set of coordinated, sequential cellular events (Nowakowski and Hayes, 1999). For instance, pregnancy is a state of sequential physiologic changes, such that an infant may be particularly susceptible to exposure during a certain developmental stage, which we call a “critical exposure window” or “window of susceptibility”. Knowledge about time windows of susceptibility could provide insight into the biologic mechanism underlying the health effect, and therefore inform prevention and treatment strategies. Metal mixture exposures may be especially harmful during prenatal and early life periods. Several metals cross the placental barrier, potentially causing injury to the fetal brain (Needham *and others*, 2011; Yoon *and others*, 2009). A previous study reported that the interaction of lead and cadmium may depend on the stage of pregnancy (Kim *and others*, 2013). In such cases, measuring exposure either in the wrong critical window or averaging exposure over the entire pregnancy when only a specific window is most relevant is a form of exposure misclassification.

There is a paucity of statistical methods to simultaneously accommodate the complex exposure–response relationship between metal mixtures and neurodevelopment while analysing data on critical windows of these exposures. Traditionally, these two research questions—(1) assessing health effects of complex mixtures and (2) identifying time windows of susceptibility—have been studied separately. Methods to address complex exposure–response relationships include principal components analysis, sparse partial least squares, classification and regression trees, random forest, cluster analysis, non-parametric Bayesian shrinkage, Bayesian mixture modeling, and weighted quantile sum regression (Chun and Keles, 2010; Billionnet *and others*, 2012; Herring, 2010; de Vocht *and others*, 2012; Diez *and others*, 2012; Roberts and Martin, 2006; Gennings *and others*, 2013). Bobb *and others* (2015) developed Bayesian kernel machine regression (BKMR) for estimating health effects of complex mixtures and conducting variable selection for exposures measured at a single time point. Meanwhile, methods for identifying time windows of susceptibility have focused on the effects of a single pollutant, such as lead (Hu *and others*, 2006; Sanchez *and others*, 2011). The models mainly use single pollutant distributed lag models to study the effect of a single toxicant assuming no interaction between time windows (Hsu *and others*, 2015; Warren *and others*, 2012, 2013; Darrow *and others*, 2011). One exception is the work of Heaton and Peng (2013), who developed a higher degree distributed lag model to account for interaction among the same exposure measured at multiple time points. However, this model still only considers exposure

to a single pollutant. Another exception is the work of [Bello and others \(2017\)](#), who developed lagged weighted quantile sum regression and tree-based distributed lag modeling for time-varying chemical mixtures. However, these methods cannot characterize the complex exposure response surface; instead, they only quantify the magnitude, not directionality, of the mixture effect, and cannot identify which mixture components are contributing positively or negatively to the mixture. To our knowledge, there are limited methods to identify critical exposure windows of multi-pollutant mixtures.

To address this gap in the statistical literature, we develop methodology to investigate how exposures to environmental mixtures during early childhood affect long-term cognitive function, and to identify specific critical windows of exposure. We introduce a new method, called lagged kernel machine regression (LKMR), to estimate the health effects of time-varying exposures to environmental mixtures, and identify critical exposure windows of a mixture. We adopt a Bayesian paradigm for inference of LKMR. We use the kernel machine regression (KMR) framework, which is popular in the statistical genetics literature where it is used primarily to test the significance of gene sets and predict risk for health outcomes ([Cai and others, 2011](#); [Maity and Lin, 2011](#)). BKMR has also been shown to effectively estimate complex exposure–response functions associated with metal mixtures ([Bobb and others, 2015](#)). We develop LKMR to handle time-varying mixture exposures. By incorporating methods from the single time point BKMR and the single exposure distributed lag model, LKMR estimates non-linear and non-additive effects of exposure mixtures while assuming these effects vary smoothly over time. To accomplish these goals, we develop a novel Bayesian penalization scheme that combines the group and fused lasso ([Yuan and Lin, 2006](#); [Tibshirani and Saunders, 2005](#)). The group lasso regularizes the exposure–response function at each time point, whereas the fused lasso shrinks the exposure–response functions from timepoints close in time towards one another. Notably, we show this can be achieved by embedding the kernel matrix into the penalty term for the grouped lasso component. We implement the method using Bayesian lasso methods ([Yuan and Lin, 2006](#); [Kyung and others, 2010](#)). Although Bayesian grouped lasso and fused lasso have been used individually, our new method combines these penalization schemes together with kernel machine methods, resulting in a novel model formulation.

LKMR can be readily used to analyse a variety of exposure–response relationships. For the purposes of this article, we focus on time-varying exposures to heavy metal mixtures and their effects on neurodevelopmental outcomes. We apply this model to data from the ongoing Early Life Exposures in Mexico and Neurotoxicology (ELEMENT) study. ELEMENT is a prospective birth cohort study followed in Mexico City with extensive neurophenotyping and covariate data collected longitudinally. In particular, a novel tooth biomarker ([Arora and others, 2014](#); [Arora and Austin, 2014](#)) was developed for application in this cohort. Tooth dentine captures exposure to multiple metals including barium (Ba), chromium (Cr), lithium (Li), manganese (Mn), and zinc (Zn) from the second trimester of pregnancy to early childhood. To our knowledge, this is the only method in which exposure data can be captured repeatedly and longitudinally over such a long time span. In our data application, we estimate the neurodevelopmental effects of joint exposures to five metals (barium, chromium, lithium, manganese, and zinc) at three exposure windows of early development (second and third trimesters of pregnancy and months 0-3 after birth).

## 2. METHODS

### 2.1. Kernel machine regression

We first review KMR as a framework for estimating the effect of a complex mixture when exposure is measured only at a single time point. We describe the model for a continuous, normally distributed outcome. For each subject  $i = 1, \dots, n$ , KMR relates the health outcome ( $y_i$ ) to  $M$  components of the exposure mixture  $\mathbf{z}_i = (z_{1i}, \dots, z_{Mi})^T$  through a flexible function,  $h(\cdot)$ , while controlling for  $C$  relevant

covariates  $\mathbf{x}_i = (x_{1i}, \dots, x_{Ci})^T$ . The model is

$$y_i = h(z_{1i}, \dots, z_{Mi}) + \mathbf{x}_i^T \boldsymbol{\beta} + \epsilon_i, \quad (2.1)$$

where  $\boldsymbol{\beta}$  represents the effects of the potential confounders, and  $\epsilon_i \stackrel{iid}{\sim} N(0, \sigma^2)$ .  $h(\cdot)$  can be estimated parametrically or non-parametrically. We employ a kernel representation for  $h(\cdot)$  in order to accommodate the possibly complex exposure–response relationship.

The unknown function  $h(\cdot)$  can be specified through basic functions or through a positive definite kernel function  $K(\cdot, \cdot)$ . Under regularity conditions, Mercer’s theorem (Cristianini and Shawe-Taylor, 2000) showed that the kernel function  $K(\cdot, \cdot)$  implicitly specifies a unique function space,  $H_K$ , that is spanned by a set of orthogonal basis functions. Thus, any function  $h(\cdot) \in H_K$  can be represented through either a set of basis functions under the primal representation, or through a kernel function under the dual representation. The kernel function uses a similarity metric  $K(\cdot, \cdot)$  to quantify the distance between the exposure profiles  $\mathbf{z}_i$  and  $\mathbf{z}_j$  for subjects  $i$  and  $j$  in the study. For example, the Gaussian kernel quantifies similarity through the Euclidean distance; the polynomial kernel, through the inner product. By specifying different kernels, one is able to control the complexity and form of the exposure–response function.

Liu and others (2007) developed least-squares kernel machine semi-parametric regression for studying genetic pathway effects. They make the connection between kernel machine methods and linear mixed models, demonstrating that (2.1) can be expressed as the mixed model

$$y_i \sim N(h_i + \mathbf{x}_i^T \boldsymbol{\beta}, \sigma^2) \quad (2.2)$$

$$\mathbf{h} = (h_1, \dots, h_n)^T \sim GP[\mathbf{0}, \tau \mathbf{K}(\cdot, \cdot)], \quad (2.3)$$

where  $\mathbf{K}$  is a kernel matrix with  $i, j$ -th element  $K(\mathbf{z}_i, \mathbf{z}_j)$ , and GP stands for Gaussian process.

## 2.2. Lagged kernel machine regression

Now assume exposures to a complex mixture are measured at multiple time points,  $t = 1, \dots, T$ , with the goal of identifying critical windows of exposure, such that we have data on the multi-pollutant exposures  $\mathbf{z}_{i,t} = (z_{1i,t}, \dots, z_{Mi,t})^T$ . We define LKMR as

$$y_i = \beta_0 + \sum_t h_t(z_{1i,t}, \dots, z_{Mi,t}) + \mathbf{x}_i^T \boldsymbol{\beta} + \epsilon_i. \quad (2.4)$$

Model (2.4) represents a multiple kernel learning model. The function  $h_t(\mathbf{z}_{i,t})$  represents the time-specific effects of exposure mixtures, while controlling for exposure at the other time windows. We use the mixed model representation proposed by Liu and others (2007), yielding

$$y_i = \beta_0 + \sum_t h_{i,t} + \mathbf{x}_i^T \boldsymbol{\beta} + \epsilon_i, \quad (2.5)$$

where  $\mathbf{h}_t = (h_{1,t}, \dots, h_{n,t})^T$  represents the (potentially complex) exposure–response random effects for the exposures  $\mathbf{z}_{1,t}, \dots, \mathbf{z}_{n,t}$  measured at time  $t$ , controlling for exposures at all other time points. It is typically the case in environmental data that exposures are highly correlated across multiple time points. Thus, naively estimating the time-specific exposure–response function without addressing this correlation can lead to unstable estimates of the  $h_t(\mathbf{z}_{i,t})$ ,  $t = 1, \dots, T$ . Accordingly, in LKMR, we impose penalization through a novel combination of Bayesian group lasso

(Yuan and Lin, 2006) and Bayesian fused lasso (Tibshirani and Saunders, 2005). The group lasso component regularizes each  $\mathbf{h}_t$  individually, as exposures at individual time windows can share similarities, and also provides a framework for incorporating the kernel matrix. The fused lasso component shrinks differences in elements of  $\mathbf{h}_t$  across neighboring time windows, and smooths adjacent  $\mathbf{h}_t$  estimates towards one another. To account for the possibility of complex non-linear and non-additive exposure–response functions at each time  $t$ , we use a novel parameterization by incorporating kernel distance functions within the group lasso implementation.

Specifically, let  $\mathbf{h} = (\mathbf{h}_1^T, \dots, \mathbf{h}_T^T)^T$ . The conditional prior of  $\mathbf{h}|\sigma^2$ , which reflects the Bayesian grouped, fused lasso penalization, is

$$\pi(\mathbf{h}|\sigma^2, \lambda_1, \lambda_2) \propto \exp \left[ \frac{-\lambda_1}{\sigma} \sum_{t=1}^T \|\mathbf{h}_t\|_{\mathbf{G}_t} - \frac{\lambda_2}{\sigma} \sum_{t=1}^{T-1} |\mathbf{h}_{t+1} - \mathbf{h}_t|_1 \right], \quad (2.6)$$

where  $\|\mathbf{h}_t\|_{\mathbf{G}_t} = (\mathbf{h}_t^T \mathbf{G}_t \mathbf{h}_t)^{1/2}$ , and corresponds to the group lasso component. We define  $\mathbf{G}_t = \mathbf{K}_t^{-1}$ , where  $\mathbf{K}_t$  denotes the kernel matrix for time  $t$  with  $i, j$  element  $K_t(\mathbf{z}_{i,t}, \mathbf{z}_{j,t})$ . Depending on a particular application, one can choose one of many different kernel functions. Meanwhile,  $|\mathbf{h}_{t+1} - \mathbf{h}_t|_1$  corresponds to the fused lasso component, where  $|\cdot|_1$  is the  $L_1$  norm of a vector.

An advantage of the model is that we can formally specify it in a hierarchical fashion, which allows for a Gibbs sampler implementation. We introduce the latent parameters  $\boldsymbol{\tau} = (\tau_1^2, \dots, \tau_T^2)$  corresponding to group lasso and  $\boldsymbol{\omega} = (\omega_1^2, \dots, \omega_{T-1}^2)$  corresponding to fused lasso which parameterize  $\Sigma_h$ , and specify the LKMR model using the following hierarchical model formulation:

$$\mathbf{y}|\mathbf{h}, \mathbf{X}, \boldsymbol{\beta}, \sigma^2 \sim N_n \left( \mathbf{X}\boldsymbol{\beta} + \sum_t \mathbf{h}_t, \sigma^2 \mathbf{I}_n \right) \quad (2.7)$$

$$\mathbf{h}|\tau_1^2, \dots, \tau_T^2, \omega_1^2, \dots, \omega_{T-1}^2, \sigma^2 \sim N_{nT} (\mathbf{0}, \sigma^2 \Sigma_h) \quad (2.8)$$

$$\tau_1^2, \dots, \tau_T^2 \stackrel{iid}{\sim} \text{gamma} \left( \frac{n+1}{2}, \sum_t \frac{\lambda_1^2}{2} \right) \quad (2.9)$$

$$\omega_1^2, \dots, \omega_{T-1}^2 \stackrel{iid}{\sim} \text{gamma} \left( 1, \frac{\lambda_2^2}{2} \right), \quad (2.10)$$

where  $\tau_1^2, \dots, \tau_T^2, \omega_1^2, \dots, \omega_{T-1}^2, \sigma^2$  are mutually independent. The form of  $\Sigma_h^{-1}$  follows from representing the Laplace (double exponential) conditional prior of  $\mathbf{h}|\sigma^2$  as a scale mixture of a normal distribution with an exponential mixing density (Andrews and Mallows, 1974). Section A of supplementary material available at *Biostatistics* online provides the full form of the variance covariance matrix  $\Sigma_h^{-1}$ , which is parameterized by  $\boldsymbol{\tau}^2 = (\tau_1^2, \dots, \tau_T^2)$ ,  $\boldsymbol{\omega}^2 = (\omega_1^2, \dots, \omega_{T-1}^2)$ , and the parameters in  $\mathbf{K}_t$ . The diagonal blocks of size  $n \times n$  arise due to the kernel structure placed on each  $\mathbf{h}_t$ , whereas the off-diagonals involving the  $\omega_t^2$  parameters serve to shrink random effects adjacent in time towards one another. Additional details on the prior specification, MCMC sampler and full conditional distributions for LKMR can be found in Section B of supplementary material available at *Biostatistics* online.

### 2.3. Predicting health effects at new time-varying exposure profiles

An advantage of LKMR is the ability to estimate and visualize the exposure–response surface for exposures measured at a given time, adjusted for exposures measured at other timepoints. Suppose we are interested in predicting the exposure–response function for  $q$  new metal mixture exposures at time  $t$ . The model is

fit to data from  $n$  subjects with exposure values measured at times  $t = 1, \dots, T$ . To interpret the model fit, interest focuses on predicting the exposure response surface  $h_t(\cdot)$  at these  $q$  new mixture values,  $\mathbf{z}_{1,t}^{\text{new}}, \dots, \mathbf{z}_{q,t}^{\text{new}}$ , where  $\mathbf{z}_{i,t}^{\text{new}} = (z_{i1,t}^{\text{new}}, \dots, z_{Mi,t}^{\text{new}})^T$ . Thus, we predict new random effects  $\mathbf{h}_t^{\text{new}} = (h_{1,t}^{\text{new}}, \dots, h_{q,t}^{\text{new}})^T$  corresponding to subjects with exposures  $\mathbf{z}_{1,t}^{\text{new}}, \dots, \mathbf{z}_{q,t}^{\text{new}}$  at time  $t$  given the observed data.

For the convenience of a simple structure for the covariance matrix, we define  $\tilde{\mathbf{h}}$  as a reordered  $\mathbf{h}$  vector, such that the time point of interest,  $t$ , is at the end of the vector. Thus,  $\tilde{\mathbf{h}} = (\mathbf{h}_1^T, \dots, \mathbf{h}_{t-1}^T, \mathbf{h}_{t+1}^T, \dots, \mathbf{h}_T^T, \mathbf{h}_t^T, (\mathbf{h}_t^{\text{new}})^T)^T$ . We reorder the corresponding covariance matrix and denote this reordered matrix as  $\sigma^2 \tilde{\Sigma}_h$ . The joint distribution of observed and new exposure profiles is:

$$\begin{pmatrix} \mathbf{h} \\ \mathbf{h}_{\text{new}} \end{pmatrix} \sim N \left[ \mathbf{0}, \sigma^2 \tilde{\Sigma}_h = \sigma^2 \begin{pmatrix} A & B \\ B^T & C \end{pmatrix}^{-1} = \sigma^2 \begin{pmatrix} \tilde{\Sigma}_{11} & \tilde{\Sigma}_{12} \\ \tilde{\Sigma}_{12}^T & \tilde{\Sigma}_{22} \end{pmatrix} \right], \quad (2.11)$$

where  $A$  represents the appropriately reordered inverse covariance matrix  $\tilde{\Sigma}_h^{-1}$ ,  $C$  denotes the  $q \times q$  matrix with  $(i, j)$ th element constructed using the inverse of the kernel function  $K(\mathbf{z}_{i,t}^{\text{new}}, \mathbf{z}_{j,t}^{\text{new}})$  evaluated at the new mixture values at the time point of interest, and  $B$  is a  $nT \times q$  matrix with  $(i, j)$ th element involving the observed mixture  $\mathbf{z}_{i,t}$  and a new mixture  $\mathbf{z}_{j,t}^{\text{new}}$  at the same time point of interest. Here,  $B$  is constructed using the inverse of the kernel function  $K(\mathbf{z}_{i,t}, \mathbf{z}_{j,t}^{\text{new}})$  and is zero otherwise. Thus,  $\tilde{\Sigma}_{11}$  is a  $nT \times nT$  matrix,  $\tilde{\Sigma}_{22}$  is a  $q \times q$  matrix, and  $\tilde{\Sigma}_{12}$  is a  $nT \times q$  matrix.

It follows that the conditional posterior distribution of  $\mathbf{h}_{\text{new}}$  is:

$$\begin{aligned} \mathbf{h}^{\text{new}} | \boldsymbol{\beta}, \boldsymbol{\tau}^2, \boldsymbol{\omega}^2, \sigma^2 &\sim N_q \left[ \tilde{\Sigma}_{12}^T \tilde{\Sigma}_{11}^{-1} \left\{ \mathbf{W}^T \mathbf{W} + \tilde{\Sigma}_{11}^{-1} \right\}^{-1} \mathbf{W}^T (\mathbf{Y} - \mathbf{X}\boldsymbol{\beta}), \right. \\ &\left. \sigma^2 \tilde{\Sigma}_{12}^T \tilde{\Sigma}_{11}^{-1} \left\{ \mathbf{W}^T \mathbf{W} + \tilde{\Sigma}_{11}^{-1} \right\}^{-1} \tilde{\Sigma}_{11}^{-1} \tilde{\Sigma}_{12} + \sigma^2 \tilde{\Sigma}_{22} - \sigma^2 \tilde{\Sigma}_{12}^T \tilde{\Sigma}_{11}^{-1} \tilde{\Sigma}_{12} \right]. \end{aligned} \quad (2.12)$$

In order to reduce computation time, we approximate the posterior mean and variance of  $\mathbf{h}_{\text{new}}$  based on the estimated posterior mean of the other parameters, specifically  $\boldsymbol{\beta}$ ,  $\boldsymbol{\tau}^2$ ,  $\boldsymbol{\omega}^2$ ,  $\sigma^2$ ,  $\lambda_1^2$  and  $\lambda_2^2$ .

### 3. SIMULATION STUDY

We conducted a simulation study to evaluate the performance of the proposed LKMR model for estimating the exposure–response function at critical exposure windows of environmental mixtures. We compared the results from LKMR to those from multiple BKMR models applied using exposure data from each window separately, as well as to those from a joint kernel BKMR (JKBKMR) which includes all chemical concentrations from all timepoints in a single kernel function.

LKMR simultaneously accounts for all time-varying exposures within a single model and allows for non-linearity and interaction between components of individual time windows, and provides estimated time-specific effects. Meanwhile, BKMR is designed for cross-sectional studies, and thus a separate BKMR model needs to be run for each time window in order to obtain time-specific effects. While BKMR does consider exposure to a multi-pollutant mixtures and allows for non-linearity and interaction, it only uses exposures at a single time point for each model.

Lastly, JKBKMR is equivalent to applying BKMR while including all time-varying exposures within a single kernel. JKBKMR allows for non-linearity and interaction among all time-specific exposures, including both within-time and between-time interactions. We note that when applying JKBKMR, the model does not readily provide estimated time-specific effects of a mixture. This is because BKMR

was designed for cross-sectional studies; therefore, JKBKMR provides a single  $\widehat{\mathbf{h}}$  for each individual, regardless of the number of time windows studied. In order to gain information on time-specific effects in JKBKMR, we use the prediction method proposed in BKMR to estimate a time-specific health effect for each individual. Specifically, we first run JKBKMR; then, at each time window of interest, we hold the metal exposures at all other time windows at their median, and use the observed metal exposures of the time point of interest to predict the time-specific  $\widehat{\mathbf{h}}$ . The true time-specific  $\mathbf{h}$  remains the same for LKMR, BKMR and JKBKMR, allowing us to make comparisons between the three methods.

Our simulation study considered a five-toxicant scenario: two toxicants (out of five) exert a gradual non-additive, non-linear effect over four time windows that are representative of pregnancy and early life. We used the following model:  $y_i = \mathbf{x}_i^T \boldsymbol{\beta} + \sum_t h_t(\mathbf{z}_{i,t}) + e_i$ , where  $\mathbf{z}_{i,t} = (z_{1i,t}, z_{2i,t}, z_{3i,t})^T$ ,  $e_i \sim N(0, 1)$ ,  $\mathbf{x}_i = (x_{1i}, x_{2i})^T$  and  $x_{1i} \sim N(10, 1)$  and  $x_{2i} \sim \text{Bernoulli}(0.5)$ . We simulated auto-correlation within toxicant exposures  $\mathbf{z}_m = (\mathbf{z}_{m,1}^T, \mathbf{z}_{m,2}^T, \mathbf{z}_{m,3}^T, \mathbf{z}_{m,4}^T)^T$  for toxicant  $m = 1, 2, 3, 4, 5$  across time, and correlation among toxicants, using the Kronecker product for the exposure correlation matrix. Four choices for auto-correlation (AR-1) within toxicants were considered: high (0.8), medium (0.5), low (0.2) and none (0). The shape of the exposure-response function  $h_t(\mathbf{z}_{i,t})$ , which was assumed to be the same at each time point, was simulated as quadratic with two-way interactions. We assumed there is no effect of exposure to the mixture at time  $t = 1$ , and a gradual increase in the effect over time, by defining  $h_t(\mathbf{z}) = \alpha_t h(\mathbf{z})$ , where  $\alpha = (\alpha_1, \alpha_2, \alpha_3, \alpha_4) = (0, 0.5, 0.8, 1.0)$  and  $h(\mathbf{z}) = z_1^2 - z_2^2 + 0.5z_1z_2 + z_1 + z_2$ .

Table 1 presents the results of this simulation study. For LKMR, BKMR, and JKBKMR, we used the quadratic kernel function, such that  $K(\mathbf{z}, \mathbf{z}') = (\mathbf{z}^T \mathbf{z}' + 1)^2$ . For each simulated data set, to assess the performance of the model for the purposes of estimating the time-specific exposure-response function, we regressed the predicted  $\widehat{\mathbf{h}}$  on  $\mathbf{h}$  for each time point. We present the intercept, slope and  $R^2$  of the regressions over 100 simulations, each with a dataset of 100 subjects. Good estimation performance occurs when the intercept is close to zero, and the slope and  $R^2$  are both close to one. We also present the root mean squared error (RMSE) and the coverage (the proportion of times the true  $h_{i,t}$  is contained in the posterior credible interval).

The results in Table 1 suggest that the LKMR significantly outperforms BKMR when there is high auto-correlation for individual mixture components across time. Specifically, as compared to BKMR, LKMR provides reductions in RMSE on the order of 70–80% when the auto-correlation in a given exposure is 0.8, and by 16–41% when the auto-correlation is 0. Credible interval coverage for LKMR is consistently at or above 95%, as compared to credible interval coverage for BKMR which is 60–74% for high (0.8) exposure auto-correlation and 82–93% for medium (0.5) auto-correlation. Furthermore, as demonstrated by a slope of  $\widehat{\mathbf{h}}$  on  $\mathbf{h}$  greater than one, BKMR estimates of  $\widehat{\mathbf{h}}$  at a given time point tend to be biased when the exposure auto-correlation is high, whereas estimates from the LKMR model are approximately unbiased under all auto-correlation scenarios. Collectively, these results demonstrate that naive application of BKMR in this setting suffers from the fact that it estimates the association between exposure at a given time but does not control for exposure at other time points. When auto-correlation in exposure among multiple exposure times is high, this lack of adjustment leads to biased estimates of an exposure effect at the time of interest, whereas when the exposures are approximately uncorrelated, there is less potential for confounding by exposure at different times. In contrast, because LKMR uses penalization to borrow information from neighboring time windows, it performs well under both high and low AR-1 scenarios, and is capable of handling time-varying mixture exposures. This simulation study demonstrates that LKMR is less biased than BKMR at estimating the exposure-response function across multiple time points.

Furthermore, the results in Table 1 suggest that the LKMR also outperforms JKBKMR. As compared to JKBKMR, LKMR provides reductions in RMSE on the order of 15–52% when auto-correlation in a given



Table 1. Simulation: Regression of  $\hat{h}$  on  $h$ . Performance of estimated  $h_t(\mathbf{z}_t)$  across 100 simulated datasets, each with  $N = 100$ . AR-1 denotes the autocorrelation in simulated data, ranging from none (0) to high (0.8). RMSE denotes the root mean squared error of the  $\hat{h}$  as compared to  $h$ . Coverage denotes the proportion of times that the true  $h$  falls within the posterior credible interval of each time point. At time window 1, there is no effect; thus, slope and  $R^2$  are not applicable to the regression of  $\hat{h}$  on  $h$ . BKMR refers to applying BKMR to exposure data from each time point separately. JKBMKR refers to applying BKMR simultaneously to data from all time points

$h$ func.	Time	Intercept	Slope	$R^2$	RMSE	Coverage	$h$ func.	Time	Intercept	Slope	$R^2$	RMSE	Coverage
LKMR	1	0.07	N/A	N/A	0.66	0.98	LKMR	1	0.05	N/A	N/A	0.62	0.98
	2	0.03	0.78	0.75	0.64	0.96		2	0.01	0.84	0.79	0.59	0.97
	3	0.01	0.83	0.89	0.69	0.95		3	-0.01	0.90	0.91	0.64	0.97
	4	-0.05	0.91	0.92	0.73	0.98		4	-0.01	0.95	0.93	0.68	0.98
BKMR	1	0.29	N/A	N/A	0.79	0.96	BKMR	1	0.21	N/A	N/A	1.18	0.93
	2	0.26	0.58	0.48	1.04	0.94		2	0.14	1.36	0.64	1.48	0.90
	3	0.24	0.72	0.71	1.16	0.92		3	0.11	1.48	0.84	1.65	0.82
	4	0.18	0.82	0.84	1.06	0.93		4	0.13	1.21	0.86	1.38	0.89
JKBMKR	1	0.01	N/A	N/A	0.72	0.98	JKBMKR	1	-0.04	N/A	N/A	0.73	0.97
	2	0.02	0.64	0.60	0.89	0.95		2	-0.04	0.64	0.57	0.94	0.95
	3	0.06	0.65	0.74	1.11	0.91		3	-0.01	0.66	0.71	1.17	0.91
	4	0.07	0.64	0.78	1.30	0.87		4	0.01	0.66	0.79	1.27	0.88
LKMR	1	0.05	N/A	N/A	0.61	0.98	LKMR	1	0.05	N/A	N/A	0.66	0.97
	2	0.02	0.82	0.77	0.62	0.97		2	0.01	0.83	0.82	0.54	0.99
	3	0.00	0.85	0.89	0.68	0.95		3	0.01	0.92	0.92	0.58	0.99
	4	-0.02	0.94	0.93	0.71	0.98		4	-0.03	0.97	0.94	0.65	0.98
BKMR	1	0.22	N/A	N/A	0.79	0.98	BKMR	1	0.19	N/A	N/A	2.67	0.74
	2	0.17	0.79	0.52	1.10	0.95		2	0.08	2.72	0.84	2.66	0.72
	3	0.16	0.98	0.75	1.18	0.93		3	0.09	2.21	0.94	2.66	0.60
	4	0.14	0.96	0.84	1.08	0.94		4	0.05	1.68	0.93	2.11	0.72
JKBMKR	1	0.01	N/A	N/A	0.70	0.98	JKBMKR	1	-0.03	N/A	N/A	0.78	0.96
	2	0.05	0.67	0.61	0.90	0.95		2	-0.02	0.53	0.41	1.12	0.94
	3	0.04	0.68	0.74	1.10	0.91		3	-0.03	0.60	0.64	1.30	0.91
	4	0.01	0.64	0.79	1.27	0.87		4	0.03	0.64	0.74	1.39	0.86

exposure is 0.8, and by 8–44% when auto-correlation is 0. Across all time windows and all levels of auto-correlation, both slope and coverage are consistently higher for LKMR than for JKBKMR. This suggests there are limitations to capturing the time-specific effects of the mixture when all time-varying exposure data is included in the same model without penalizing within time windows or across time windows. Taken together, the simulations suggest that across a range of auto-correlation scenarios, LKMR outperforms both BKMR and JKBKMR.

To further explore the ability of LKMR to estimate time-specific effects of the exposure mixture, we consider several other simulation scenarios in Section C of supplementary material available at *Biostatistics* online. Table S1 of supplementary material available at *Biostatistics* online considers a different exposure response function than used in Table 1 in order to study the performance of LKMR when the shape of the exposure-response surface differs at different time windows. Table S2 of supplementary material available at *Biostatistics* online considers the performance of LKMR for a larger number of toxicants (10 toxicants, 4 time windows), while Table S3 of supplementary material available at *Biostatistics* online considers the performance of LKMR for a larger number of time windows (10 time windows, 5 toxicants). Lastly, Table S4 supplementary material available at *Biostatistics* online assess the performance of LKMR for a larger number of toxicants and time windows (10 toxicants, 10 time windows). In general, LKMR has a lower RMSE and a higher  $R^2$  than BKMR or JKBKMR. Across all three methods, when a larger number of time windows and/or number of toxicants is studied while holding the sample size constant, performance generally decreases. However, LKMR is still better able to estimate the time-specific exposure–response surface, even if its effects are sometimes attenuated.

#### 4. APPLICATION TO ELEMENT DATA SET

We applied LKMR to analyse the association between neurodevelopment and metal mixture exposures in the ELEMENT study conducted in Mexico City. In a pilot study ( $n = 84$ ) nested within this larger cohort study, we estimated associations between exposure to metal mixtures both before and after birth and the visual spatial subtest score of the Wide Range Assessment of Visual Motor Abilities (WRAVMA) (Adams and Sheslow, 1995) administered at approximately 8 years of age. Concentrations of metals barium, cadmium, lithium, manganese, and zinc were measured in tooth dentine, which provides time-specific measures of exposure over both the pre- and post-natal period for each child (Arora and Austin, 2014). These time-varying exposures were averaged to reflect three biologically relevant time windows: second and third trimesters of pregnancy, and months 0–3 after birth (early life). Auto-correlation varied for different metals, ranging from the lowest value of 0.56 for manganese to the highest value of 0.93 for lithium. We controlled for child gender, maternal IQ, and child hemoglobin at year two. In conducting the analysis, exposure covariates were logged, centered, and scaled. Confounder variables were also centered and scaled.

As a preliminary analysis, we first fit a linear regression model using all five metal exposures from all three time windows to identify exposures and time point(s) of significance, presented in Figure S2 of supplementary material available at *Biostatistics* online. The model included only main effects of each metal at each time window. Second trimester zinc had a positive association with neurodevelopment ( $p = 0.04$ ), and third trimester zinc had a negative association with neurodevelopment ( $p = 0.006$ ). Third trimester manganese was positively associated with neurodevelopment ( $p = 0.04$ ), and early life manganese was negatively associated with neurodevelopment ( $p = 0.006$ ). Second trimester chromium had a positive association ( $p = 0.02$ ), and early life chromium had a negative association ( $p = 0.02$ ). These results suggest that under assumptions of a linear regression model, there is some evidence of an exposure-response relationship across multiple time points, which warrants further exploration using LKMR.

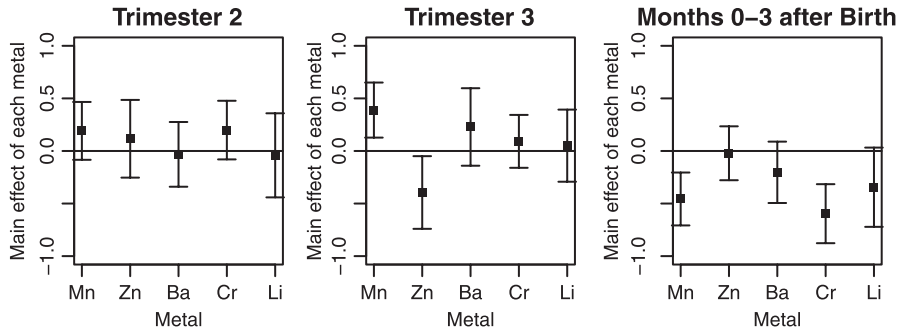


Fig. 1. LKMR estimated relative importance of each metal at three critical windows for ELEMENT data. Plot of the estimated relative importance of each metal, as quantified by the difference in the mean response at the 75th percentile versus the 25th percentile of a given metal exposure, while holding all other metal exposures constant at their median values.

We then applied LKMR to study time-varying metal mixture exposure effects during early life on visual spatial ability. Previous literature has shown an inverted- $u$  relationship between metals and neurodevelopment (Claus Henn *and others*, 2010); thus, we chose  $\mathbf{K}_t$  to be a quadratic kernel, such that  $K(\mathbf{z}, \mathbf{z}') = (\mathbf{z}^T \mathbf{z}' + 1)^2$ . The quadratic kernel allows for the possibility that a given metal can have different effects (beneficial and toxic) at different exposure levels. Due to the moderate sample size, we also assessed the sensitivity of LKMR to different choices of prior hyperparameters, and found that the results to be stable, and would recommend sensitivity analysis when applying LKMR to any given application. Using LKMR, we estimated the relative importance of each metal separately for each time window. Relative importance is quantified as the difference in the estimated exposure-response function when a single metal is at a high exposure level (75th percentile) versus a low exposure level (25th percentile), holding all other metals constant at their median exposure levels. Figure 1 shows the results, which are similar to those from simple linear regression. Third trimester zinc was detected to be negatively associated with neurodevelopment. The results also suggest evidence of a positive association of third trimester manganese with neurodevelopment, which shifts to a negative association after birth. This qualitatively different (positive vs. negative) association between manganese and neurodevelopment for prenatal vs. postnatal manganese exposure is particularly intriguing, as manganese is both an essential nutrient and a toxicant. It could be that the developing fetus needs manganese prenatally and receives it via the mother, whereas post-natal levels reflects environmental exposures that are more harmful.

As the estimated relative importance from LKMR indicate effects of manganese and zinc, we focus on those two metals when exploring the exposure-response relationship. Because the exposure-response surface is five-dimensional, we use heat maps and cross-sectional plots to reduce dimensionality and graphically depict the exposure-response relationship. Figure 2 presents the plot of the posterior mean of the exposure-response surface of manganese and zinc at the median exposure levels of barium, cadmium, and lithium estimated using LKMR. The shape of the surface at the second trimester suggests an interaction between manganese and zinc, which will be further explored below. Also, the results suggest that the direction of the association changes at birth. At the third trimester, high manganese and moderate zinc exposures are associated with higher scores, while after birth, low manganese and a range of zinc exposures are associated with higher scores.

To further reduce dimensionality, and to show estimates of the posterior uncertainty of the exposure-response function, Figure 3 depicts the cross-section of the exposure-response surface for manganese, at low and high zinc concentrations and median values of the barium, chromium, and lithium exposures. These results suggest that the association between manganese exposure and visual spatial score depends on




Fig. 2. LKMR estimated time-specific exposure response functions applied to ELEMENT data. Plot of the estimated exposure-response surface for Mn and Zn, at the median of Ba, Cd, Li.

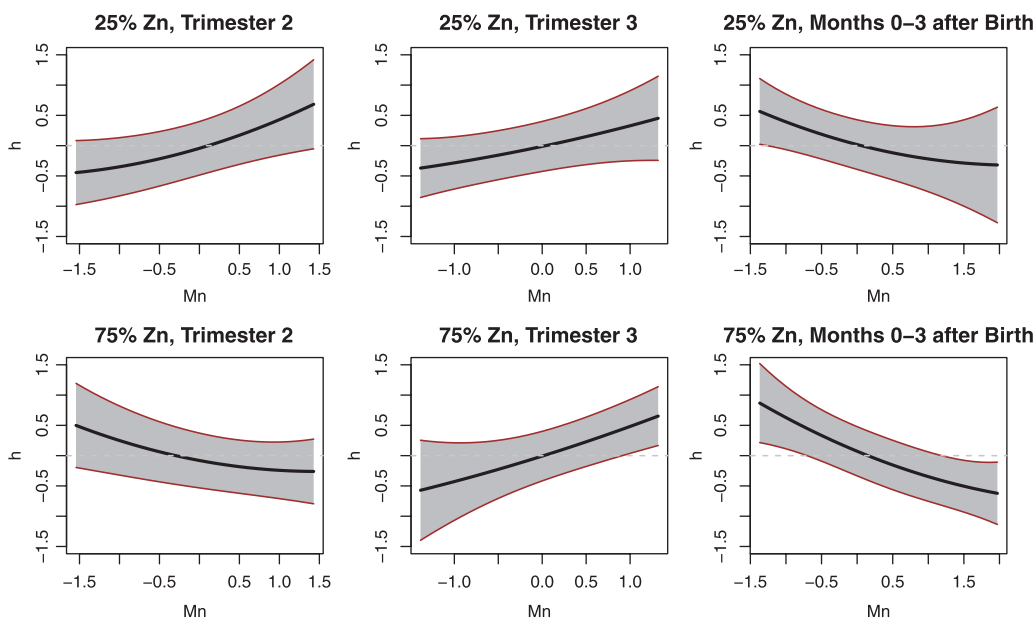


Fig. 3. LKMR estimated time-specific exposure-response functions for Mn at low and high Zn levels applied to ELEMENT data. Plot of the cross-section of the estimated exposure-response surface for Mn, at the 25th (top panel) and 75th (bottom panel) of Zn exposure, holding Ba, Cr, and Li constant at median exposures.

exposure timing. Comparing the top panel (low zinc) to the bottom panel (high zinc), we detect suggestion of a manganese-zinc interaction, specifically effect modification in the presence of higher zinc levels during the second trimester. At the second trimester, there is a positive association between manganese exposure and visual spatial score in the presence of low zinc levels. However, the association becomes negative in the presence of high zinc levels. Notably, this interaction is not suggested in the plot of the relative importance, where the effects of manganese and zinc are both non-significant at the second trimester. In the cross-sectional plot Figure 3, we also note evidence of a positive association between manganese and neurodevelopment before birth at the third trimester, and a negative one after birth. Lastly, the cross-sectional graphs suggest the effects are mainly linear, indicating that a quadratic kernel is sufficient to capture the exposure-response relationship.

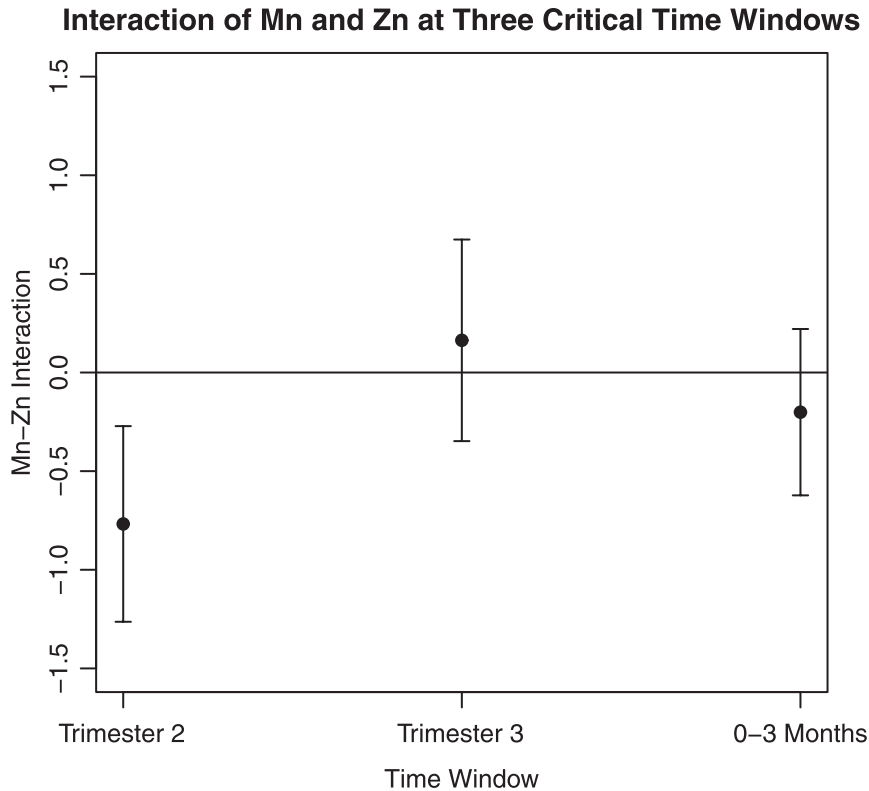


Fig. 4. LKMR estimated Mn-Zn interaction at three critical windows for ELEMENT data. Plot of the estimated interaction effect between Mn and Zn, holding Ba, Cr, and Li constant at median exposures. The interaction effect was calculated as follows. First, we estimated the exposure-response effect for high (75th) versus low (25th) manganese exposures, at high zinc levels. Next, we estimated the exposure-response effect for high versus low manganese exposures, at low zinc levels. The difference between the two estimated exposure-response effects quantifies the manganese-zinc interaction. All other metal exposures are held constant at their median values.

In Figure 4, we focus on the estimated interaction effect between manganese and zinc at the three critical time windows. First, we estimated the exposure-response effect for high (75th percentile) versus low (25th percentile) manganese exposures, at high zinc levels and median levels of barium, chromium and lithium. Next, we estimated the exposure-response effect for high versus low manganese exposures, at low zinc levels and median levels of barium, chromium and lithium. The difference between the two estimated exposure-response effects quantifies the manganese-zinc interaction. The results indicate that there is a significant manganese-zinc interaction at the second trimester.

To complete our case study, we compare the results under LKMR to those obtained by BKMR applied using data from each critical window separately, and to those obtained from JKBKMR, which captures all time-varying exposures within a single kernel, in Section D of supplementary material available at *Biostatistics* online. Figure S2 supplementary material available at *Biostatistics* online shows the relative importance of each metal as estimated by BKMR and JKBKMR. The results from BKMR suggest that when focused on manganese and zinc, only third trimester zinc exposure was significantly negatively associated with visual spatial score. This differs from the results found under LKMR and the linear model, which indicate a positive association of third trimester manganese and a negative association of third

trimester zinc. Meanwhile, the results from JKBKMR do not identify any metals at any time points as being significantly associated with the outcome. Figure S3 of supplementary material available at *Biostatistics* online depicts the posterior mean of the exposure–response surface of manganese and zinc at the median of barium, cadmium, and lithium for BKMR and JKBKMR. Neither BKMR nor JKBKMR detect associations or interactions between second trimester manganese and zinc exposures with neurodevelopment. However, under LKMR, there was suggestion of a second trimester manganese–zinc interaction. In months 0–3 after birth, however, findings from LKMR, BKMR, and JKBKMR are generally consistent, with low manganese exposure associated with higher neurodevelopmental scores across a range of zinc exposures. However, the estimated time-specific exposure response surfaces for BKMR and JKBKMR have an attenuated signal compared to that from LKMR which identifies an interaction at the second trimester. Lastly, Figure S4 of supplementary material available at *Biostatistics* online depicts the predicted cross-sectional plot for BKMR and JKBKMR, which do not suggest a Mn–Zn interaction any any time window, in contrast to LKMR. Taken together, these findings further suggest that both BKMR and JKBKMR have limited ability to detect a signal, which may be due to serial correlation and confounding by exposure at the other time points.

## 5. DISCUSSION

In this article, we have developed a LKMR model that uses Bayesian regularization to analyse data on time-varying exposures of environmental mixtures to identify critical windows of exposure in children’s health. A unique exposure biomarker that captures the temporal variation in exposure in short time windows makes this methodology possible. The kernel framework allows for a flexible specification of the unknown exposure–response relationship. We use a Bayesian formulation of the group lasso, which regularizes each kernel surface, and the fused lasso, which smooths multivariate exposure–response surfaces over time. Our method can account for auto-correlation of mixture components over time while exploring the possibility of non-linear and non-additive effects of individual exposures. A key contribution of this article is the incorporation of the kernel machine framework into a grouped, fused lasso framework.

We demonstrated that the LKMR method achieves large gains over a simpler cross-sectional approach that considers each critical window separately, particularly when serial correlation among the time-varying exposures is high. LKMR also outperforms a method that accounts for all time-varying exposure data within a single kernel and does not adjust for time-specific effects. We applied LKMR to analyse associations between neurodevelopment and metal mixtures in the ELEMENT cohort. In the presence of complex exposure–response relationships that can vary with the timing of exposures, LKMR is a promising method to quantify health effects and identify time windows of susceptibility. In the application of LKMR to data from the ELEMENT study, LKMR, which uses information from neighboring time windows through penalization, is able to detect effect modification that was missed by both BKMR and JKBKMR. Specifically, LKMR detected an interesting interaction between manganese and zinc. At low levels of zinc, manganese exposure at the second trimester of pregnancy is positively associated with neurodevelopment. However, this positive association shifts after birth, at which point manganese is negatively associated with visual spatial ability. This suggests manganese functions as a trace element and an essential nutrient before birth, and as a toxicant after birth. Furthermore, this effect is not present under high exposure levels of zinc at the second trimester. The finely detailed interaction effect is captured by LKMR but not by BKMR or JKBKMR. We also note that in this case study, we saw similarities between results from LKMR and those from the linear model. The linear model, which is used for preliminary analysis, makes strong assumptions of linearity and additivity. It is important to use a flexible method like LKMR to check these assumptions as they can be violated; for example, in the case study from BKMR, the authors detected a non-linear and non-additive exposure–response function for two metals (Bobb and others, 2015). LKMR also provides added advantage over the linear model by detailing the

interaction between Mn and Zn, through visualization of the exposure response surface from heatmaps and cross-sectional plots.

LKMR is best-suited for analysis of environmental mixture data where exposures are measured at a small number of time windows, as may arise from exposures measured in blood or other biomarkers. Alternative data structures, such as exposures measured semi-continuously over time—for example, air pollution measurements measured weekly—may require alternative methods based on functional data analysis (Wang and others, 2015; Malloy and others, 2010) or distributed lag modeling (Bello and others, 2017).

There is currently limited knowledge on the health effects of exposure to chemical mixtures. Therefore, the development of statistical methods that can handle the complexity of multi-pollutant mixtures, whose effects may vary over time, can make an important contribution to environmental health research.

#### SUPPLEMENTARY MATERIAL

Supplementary material is available at <http://biostatistics.oxfordjournals.org>.

#### ACKNOWLEDGMENTS

The authors thank the associate editor and two anonymous reviewers for their invaluable comments and suggestions. The computations in this paper were run on the Odyssey cluster supported by the FAS Division of Science, Research Computing Group at Harvard University. *Conflict of Interest*: None declared.

#### FUNDING

This work was supported by the National Institutes of Health [ES000002, ES007142, ES022986, ES024332-01A1, ES025453, ES026033, OD023286, and OD023337]; and the National Science Foundation [1514970 and 1614102].

#### REFERENCES

- ADAMS, W. AND SHESLOW, D. (1995). *Wide Range Assessment of Visual Motor Abilities*. Wilmington, DE: Wide Range.
- ANDREWS, D. F. AND MALLOWS, C. L. (1974). Scale mixtures of normal distributions. *Journal of the Royal Statistical Society, Series B (Methodological)* **36**, 99–102.
- ARORA, M. AND AUSTIN, C. (2014). Teeth as a biomarker of past chemical exposure. *Current Opinion in Pediatrics* **25**, 261–267.
- ARORA, M., AUSTIN, C., SARRAPPOUR, B., HERNÁNDEZ-ÁVILA, M., HU, H. AND WRIGHT, R. O. (2014). Determining prenatal, early childhood and cumulative long-term lead exposure using micro-spatial deciduous dentine levels. *PLoS One* **9**, e97805.
- BELLINGER, D. C. (2008). Very low lead exposures and children's neurodevelopment. *Current Opinion in Pediatrics* **20**, 172–7.
- BELLO, G. A., AUSTIN, C., HORTON, M. K., WRIGHT, R. O. AND GENNINGS, C. (2017). Extending the distributed lag model framework to handle chemical mixtures. *Environmental Research* **156**, 253–264.
- BILLIONNET, C., SHERRILL, D. AND ANNESI-MAESANO, I. (2012). Estimating the health effects of exposure to multi-pollutant mixture. *Annals of Epidemiology* **22**, 126–41.
- BOBB, J. F., VALERI, L., CLAUS HENN, B., CHRISTIANI, D. C., WRIGHT, D. O., MAZUMDAR, M., GODLESKI, J. J. AND COULL, B. A. (2015). Bayesian kernel machine regression for estimating the health effects of multi-pollutant mixtures. *Biostatistics* **16**, 493–508.

- CAI, T., TONINI, G. AND LIN, X. (2011). Kernel machine approach to testing the significance of multiple genetic markers for risk prediction. *Biometrics* **67**, 975–86.
- CARLIN, D. J., RIDER, C. V., WOYCHICK, R. AND BIRNBAUM, L. S. (2013). Unraveling the health effects of environmental mixtures: an NIEHS priority. *Environmental Health Perspectives* **121**, A6–A8.
- CHUN, H. AND KELES, S. (2010). Sparse partial least squares regression for simultaneous dimension reduction and variable selection. *Journal of the Royal Statistical Society, Series B (Methodological)* **72**, 3–25.
- CLAUS HENN, B., COULL, B. A. AND WRIGHT, R. O. (2014). Chemical mixtures and children’s health. *Current Opinion in Pediatrics* **26**, 223–229.
- CLAUS HENN, B., ETTINGER, A. S., SCHWARTZ, J., TELLEZ-ROJO, M. M., LAMADRID-FIGUEROA, H., HERNANDEZ-AVILA, M., SCHNAAS, L., AMARASIRIWARDENA, C., BELLINGER, D. C., HU, H. and others (2010). Early postnatal blood manganese levels and children’s neurodevelopment. *Epidemiology* **21**, 433–439.
- CRISTIANINI, N. AND SHAW-ETAYLOR, J. (2000). *An Introduction to Support Vector Machines*. Cambridge: Cambridge University Press.
- DARROW, L. A., KLEIN, M., STRICKLAND, M. J., MULHOLLAND, J. A. AND TOLBERT, P. E. (2011). Ambient air pollution and birth weight in full-term infants in Atlanta, 1994–2004. *Environmental Health Perspectives* **119**, 731–737.
- DE VOCHT, F., CHERRY, N. AND WAKEFIELD, J. (2012). A Bayesian mixture modeling approach for assessing the effects of correlated exposures in case-control studies. *Journal of Exposure Science & Environmental Epidemiology* **22**, 352–60.
- DIEZ, D. M., DOMINICI, F., ZARUBIAK, D. AND LEVY, J. I. (2012). Statistical approaches for identifying air pollutant mixtures associated with aircraft departures at Los Angeles International Airport. *Environmental Science & Technology* **46**, 8229–35.
- GENNINGS, C., CARRICO, C., FACTOR-LITVAK, P., KRIGBAUM, N., CIRILLO, P. M. AND COHN, B. A. (2013). A Cohort study evaluation of maternal PCB exposure related to time to pregnancy in daughters. *Environmental Health* **12**, 66.
- HEATON, M. J. AND PENG, R. D. (2013). Extending distributed lag models to higher degrees. *Biostatistics* **15**, 398–412.
- HERRING, A. H. (2010). Nonparametric bayes shrinkage for assessing exposures to mixtures subject to limits of detection. *Epidemiology* **21**, S71–6.
- HSU, L. H., CHIU, L. H., COULL, B. A., KLOOG, I., SCHWARTZ, J., LEE, A., WRIGHT, A. AND WRIGHT, R. J. (2015). Prenatal particulate air pollution and asthma onset in urban children: identifying sensitive windows and sex differences. *American Journal of Respiratory and Critical Care Medicine* doi: 10.1164/rccm.201504-0658OC.
- HU, H., TELLEZ-ROJO, M. M., BELLINGER, D., SMITH, D., ETTINGER, A. S., LAMADRID-FIGUEROA, H., SCHWARTZ, J., SCHNAAS, L., MERCADO-GARCIA, A. AND HERNANDEZ-AVILA, M. (2006). Fetal lead exposure at each stage of pregnancy as a predictor of infant mental development. *Environmental Health Perspectives* **114**, 1730–1735.
- KIM, Y., HA, E. H., PARK, H., HA, M., KIM, Y., HONG, Y. C., KIM, E. J. AND KIM, B. N. (2013). Prenatal lead and cadmium co-exposure and infant neurodevelopment at 6 months of age: the Mothers and Children’s Environmental Health (MOCEH) study. *Neurotoxicology* **35**, 15–22.
- KYUNG, M., GILL, J., GHOSH, M. AND CASELLA, G. (2010). Penalized regression, standard errors and Bayesian lassos. *Bayesian Analysis* **5**, 369–412.
- LIU, D., LIN, X. AND GHOSH, D. (2007). Semiparametric regression of multidimensional genetic pathway data: least-squares kernel machines and linear mixed models. *Biometrics* **63**, 1079–88.
- MAITY, A. AND LIN, X. (2011). Powerful tests for detecting a gene effect in the presence of possible gene-gene interactions using garrote kernel machines. *Biometrics* **67**, 1271–84.



- MALLOY, E. J., MORRIS, J. S., ADAR, S. D., SUH, H., GOLD, D. R. AND COULL, B. A. (2010). Wavelet-based functional linear mixed models: an application to measurement error-corrected distributed lag models. *Biostatistics* **11**, 432–452.
- NEEDHAM, L. L., GRANDJEAN, P., HEINZOW, B., JORGENSEN, P. J., NIELSEN, F., PATTERSON, D. G., SJODIN, A., TURNER, W. E. AND WEIHE, P. (2011). Partition of environmental chemicals between maternal and fetal blood and tissues. *Environmental Science & Technology* **45**, 1121–6.
- NOWAKOWSKI, R. S. AND HAYES, N. L. (1999). Cns development: an overview. *Development and Psychopathology* **11**, 395–417.
- ROBERTS, S. AND MARTIN, M. A. (2006). The use of supervised principal components in assessing multiple pollutant effects. *Environmental Health Perspectives* **114**, 1877–1882.
- SANCHEZ, B. N., HU, H., LITMAN, H. J. AND TELLEZ-ROJO, M. M. (2011). Statistical methods to study timing of vulnerability with sparsely sampled data on environmental toxicants. *Environmental Health Perspectives* **119**, 409–415.
- TIBSHIRANI, R. AND SAUNDERS, M. (2005). Sparsity and smoothness via the fused lasso. *Journal of the Royal Statistical Society, Series B (Methodological)* **67**, 91–108.
- WANG, J. L., CHIOU, J. M. AND MULLER, H. G. (2015). Review of functional data analysis. *Annual Review of Statistics* **7**, 1–41.
- WARREN, J., FUENTES, M., HERRING, A. AND LANGLOIS, P. (2012). Spatial-temporal modeling of the association between air pollution exposure and preterm birth: identifying critical windows of exposure. *Biometrics* **68**, 1157–1167.
- WARREN, J., FUENTES, M., HERRING, A. AND LANGLOIS, P. (2013). Air pollution metric analysis while determining susceptible periods of pregnancy for low birth weight. *Obstetrics and Gynecology* **2013**, 1–9.
- YOON, M., NONG, A., CLEWELL, H. J., TAYLOR, M. D., DORMAN, D. C. AND ANDERSEN, M. E. (2009). Evaluating placental transfer and tissue concentrations of manganese in the pregnant rat and fetuses after inhalation exposures with a pbpk model. *Toxicological Sciences* **112**, 44–58.
- YUAN, M. AND LIN, Y. (2006). Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society, Series B (Methodological)* **68**, 49–67.

[Received December 6, 2016; revised July 7, 2017; accepted for publication July 21, 2017]