

A BAYESIAN REGRESSION TREE APPROACH TO IDENTIFY THE EFFECT OF NANOPARTICLES' PROPERTIES ON TOXICITY PROFILES¹

BY CECILE LOW-KAM, DONATELLO TELESKA, ZHAOXIA JI,
HAIYUAN ZHANG, TIAN XIA, JEFFREY I. ZINK AND ANDRE E. NEL

University of California, Los Angeles

We introduce a Bayesian multiple regression tree model to characterize relationships between physico-chemical properties of nanoparticles and their in-vitro toxicity over multiple doses and times of exposure. Unlike conventional models that rely on data summaries, our model solves the low sample size issue and avoids arbitrary loss of information by combining all measurements from a general exposure experiment across doses, times of exposure, and replicates. The proposed technique integrates Bayesian trees for modeling threshold effects and interactions, and penalized B-splines for dose- and time-response surface smoothing. The resulting posterior distribution is sampled by Markov Chain Monte Carlo. This method allows for inference on a number of quantities of potential interest to substantive nanotoxicology, such as the importance of physico-chemical properties and their marginal effect on toxicity. We illustrate the application of our method to the analysis of a library of 24 nano metal oxides.

1. Introduction. The increasing use of engineered nanomaterials (ENM) in hundreds of consumer products has recently raised concern about their potential effect on the environment and human health in particular. In nanotoxicology, in vitro dose-escalation assays describe how cell lines or simple organisms are affected by increased exposure to nanoparticles. These assays help determine hazardous materials and exposure levels. Standard dose-escalation studies are sometimes completed by more general exposure escalation protocols, where a biological outcome is measured against both increasing concentrations and durations of exposure. Cost and timing issues usually only allow for a small number of nanoparticles to be comprehensively screened in any study. Therefore, both one- and two-dimensional escalation experiments are often characterized by small sample sizes. Furthermore, data exhibits natural clusters related to varying levels of nanoparticles bio-activity. The two case studies presented in Section 6 provide an overview of the structure of typical data sets obtained with both experimental protocols.

Received February 2014; revised November 2014.

¹Supported by U.S. Public Health Service Grant U19 ES019528 (UCLA Center for Nanobiology and Predictive Toxicology) and supported by the National Science Foundation and the Environmental Protection Agency under Cooperative Agreement Number DBI-0830117.

Key words and phrases. Bayesian CART, nanotoxicology, P-splines, regression trees.

Beyond dose-response analysis, nanomaterial libraries are also designed to investigate how a range of physical and chemical properties (size, shape, composition, surface characteristics) may influence ENM's interactions with biological systems. The nano-informatics literature reports several Quantitative Structure–Activity Relationship (QSAR) models. This exercise is conceived as a framework for predictive toxicology, under the assumption that nanoparticles with similar properties are likely to have similar effects. Most of existing QSAR models summarize or integrate experimental data across times, doses and replicates as a pre-processing step, before applying traditional data mining or statistical algorithms for regression. For example, Liu et al. (2011) use a modified Student's *t*-statistic to discretize outputs in two classes (toxic or nontoxic) and a logistic regression model to relate toxicity to physico-chemical variables. Zhang et al. (2012) use the area under the dose-response curve as a global summary of toxicity and they model dependence on predictors via a regression tree. Both approaches, while reasonably sensible, ignore the uncertainty associated with data summaries and can lead to unwarranted conclusions as well as unnecessary loss of information. Patel et al. (2014) summarize toxicity profiles using a new definition of toxicity, called *the probability of toxicity*, which is defined as a linear function of nanoparticle physical and chemical properties. While this last approach solves the issue of uncertainty propagation, it still makes it impossible to predict full dose-response curves from nanoparticle characteristics. Moreover, the use of regression trees is inherently appealing, as they are able to model nonlinear effects and interactions without compromising interpretation. We aim to extend regression tree models to account for structured multivariate outcomes, defined as toxicity profiles of nanoparticles, measured over a general exposure escalation domain.

Multivariate extensions of the regression tree methodology have been proposed by Segal (1992). In this paper, the original tree-building algorithm of Breiman et al. (1984) is modified to handle multivariate responses for commonly used covariance matrices, such as independence or autoregressive structures. De'ath (2002) proposes a similar method for an independent covariance structure. Yu and Lambert (1999) develop regression tree models for functional data, by representing each individual response as a linear combination of spline basis functions and using the estimated splines coefficients in multivariate regression trees. An alternative for longitudinal responses consists of combining a tree model and a linear model: Sela and Simonoff (2012) replace the fixed effects of the traditional linear mixed effects model by a regression tree. The linear random effects are unchanged. Yu et al. (2010) fit a semi-parametric model, containing a linear part and a tree part, for multivariate outcomes in genetics. The linear part is used to model main effects of some genetic or environmental exposures. The nonparametric tree part approximates the joint effect of these exposures. Finally, Galimberti and Montanari (2002) develop regression tree models for longitudinal data with time-dependent covariates. In this setting, measures for the same individual can belong to different terminal nodes.

Other extensions of standard regression trees include Bayesian approaches, where tree parameters become random variables. Chipman, George and McCulloch (1998) introduce a Bayesian regression tree model for univariate responses. The method is based on a prior distribution and a Metropolis–Hastings algorithm which generates candidate trees and identifies the most promising ones. This methodology has since been extended to so-called *treed* models, where a parametric model is fitted in each terminal node [Chipman, George and McCulloch (2002)], to a sum-of-trees model [Chipman, George and McCulloch (2010a)], and to incorporate spatial random effects for merging data sets [Zhang, Shih and Müller (2007)]. Gramacy and Lee (2008) model nonstationary spatial data by combining Bayesian regression trees and Gaussian processes in the leaves. This approach is extended to the multivariate Gaussian process with separable covariance structure in Konomi et al. (2014).

Building on previous contributions, we propose a new method to analyze the relationship between nanoparticles physico-chemical properties and their toxicity in exposure escalation experiments. We extend the Bayesian methodology of Chipman, George and McCulloch (1998) to allow for dose- and time-response kinetics in terminal nodes. Our work is closely related to the methodology introduced in Konomi et al. (2014). However, our model is specifically adapted to exposure escalation experiments, as observations for the same nanoparticle at different doses and times cannot fall in separate leaves of the tree. Therefore, the binary splits of the tree only capture structure activity relationships instead of the general increase of toxicity with exposure.

A global covariance structure accounts for correlation between measurements at different doses and times for the same nanoparticle. Our approach is able to model nonlinear effects and potential interactions of physico-chemical properties without making parametric assumptions about toxicity profiles. It also addresses the issues associated with conventional QSAR models by combining evidence across measurements for all doses and times in a general experimental design. The proposed model is particularly versatile, as it provides scores of importance for physico-chemical properties and visual assessment of the marginal effect of these properties on toxicity.

The rest of this paper is organized as follows: Section 2 describes the regression model for dose-response data and Section 3 describes the corresponding prior model. The resulting posterior distribution and the associated MCMC algorithm are presented in Section 4. The model is extended to the case of dose- and time-response surfaces in Section 5. The method is applied to a library of 24 metal oxides in Section 6 and Section 7 concludes this paper with a discussion.

2. Regression tree formulation.

2.1. *Sampling model.* We first consider the case of a typical dose escalation experiment, where a biological outcome is measured over a protocol of increased

nanoparticle concentration. This case will be expanded in Section 5 to include more general exposure escalation designs.

Let $y_{ik}(d)$ denote a real-valued response associated with exposure to nanoparticle i and replicate k at dose d , for $i \in \{1, \dots, I\}$, $k \in \{1, \dots, K\}$ and $d \in [0, D]$. We assume that y has been appropriately normalized and purified of experimental artifacts. For the two case studies of Section 6, normalization was performed for each tray by subtracting a baseline mean response, measured in control wells where cells were not exposed to any nanoparticle. After normalization, we indeed assume independence between wells exposed to different nanomaterials on the same tray. Current experimental protocols only allow for the observation of the outcome y as it varies in association with a discrete prescription of dose-escalation. However, for notational convenience and without loss of generality, we maintain that y shall be observed for any dose level d ranging from no exposure ($d = 0$) to a maximal nanoparticle concentration level ($d = D$). Let also $\mathbf{x}_i = (x_{i1}, \dots, x_{ip})$ be a p -dimensional vector of continuous physico-chemical characteristics or predictors associated to nanoparticle i . We assume

$$(2.1) \quad y_{ik}(d) = f(\mathbf{x}_i, d) + \varepsilon_{ik}(d),$$

where f is a random mean function, depending on the dose level d and nanoparticle characteristics \mathbf{x}_i , and $\varepsilon_{ik} \sim N(0, \sigma_d^2)$. More precisely, f is defined by a regression tree \mathcal{T} on the predictor space and a functional model for dose-response curves in the terminal nodes of \mathcal{T} . Full details about the proposed mean structure are described in the following section.

Given f , we assume that outcomes are independent across nanoparticles and, for any nanoparticle i , $\text{Cov}(\varepsilon_{ik}(d), \varepsilon_{ik'}(d')) = \sigma^2 \varphi_D^{|d-d'|}$, with $\varphi_D \in [0, 1]$. In this setting, two outcomes associated with the same nanoparticle at similar doses are assumed to be more correlated than measurements taken at distant doses, for any replicate. The major advantage of this assumption is related to a reduced representation of a high-dimensional covariance matrix, which is now fully characterized in terms of a 1-dimensional variance parameter σ^2 and a 1-dimensional correlation φ_D .

2.2. Mean structure. The binary tree \mathcal{T} recursively splits the predictor space into two subspaces, according to criteria of the form $x_{.j} \leq a$ vs $x_{.j} > a$, for $a \in \mathbb{R}$ and $j \in \{1, \dots, p\}$. Each split defines two new nodes of the tree, corresponding to two newly created subspaces of predictors. Let n be the set of terminal nodes of tree \mathcal{T} .

We model the dose-response curves in each terminal node as a linear combination of spline basis functions. Unlike parametric models such as log-logistic, spline functions do not assume a particular shape for the curve. This makes our model fully applicable to sub-lethal biological assays, which are not expected to follow a sigmoidal dose-response dynamic. However, if needed, the spline model

can easily allow for possible shape constraints, such as monotonicity, by using a modified basis [Ramsay (1998)]. This flexibility makes the use of spline basis representations potentially preferable to Gaussian process priors or similar smoothers. A formal comparison is, however, outside the scope of this manuscript. Our chosen functional representation is easily extended to two-dimensional response surfaces (Section 5). Let $\mathcal{B}_1(\cdot), \dots, \mathcal{B}_{m_D+\delta}(\cdot)$ denote $m_D + \delta$ uniform B-spline basis functions of order δ on $[0, D]$, with m_D fixed knots. Following Eilers and Marx (1996), we avoid choosing the location of spline interior knots by deliberately overfitting curves with a number of knots coinciding with the dose-escalation grid. Adaptive smoothness is determined by using a penalty on adjacent coefficients, via a smoothing prior that will be presented in Section 3.

If \mathbf{x}_i is in the subset corresponding to the r th terminal node of \mathcal{T} , $f(\mathbf{x}_i, d) = \sum_{\ell=1}^{m_D+\delta} \beta_{r\ell} \mathcal{B}_\ell(d)$. We will denote with $\boldsymbol{\beta}_r = (\beta_{r1}, \dots, \beta_{r m_D+\delta})'$ the vector of splines coefficients defining the expected dose-response trajectory in the r th terminal node. Furthermore, we let $\boldsymbol{\beta}$ define the random set of spline coefficients, including $\boldsymbol{\beta}_r$ from all terminal nodes ($r = 1, \dots, n$). The Bayesian model is completed by prior distributions on \mathcal{T} , $\boldsymbol{\beta}$, σ^2 and φ_D .

3. Prior model. We first introduce the general dependence structure of the prior, before describing each parameter’s prior distribution. We follow Chipman, George and McCulloch (1998), and assume that the tree is independent of variance components σ^2 and φ_D :

$$(3.1) \quad p(\mathcal{T}, \boldsymbol{\beta}, \sigma^2, \varphi_D) = p(\mathcal{T}, \boldsymbol{\beta})p(\sigma^2)p(\varphi_D) = p(\boldsymbol{\beta}|\mathcal{T})p(\mathcal{T})p(\sigma^2)p(\varphi_D).$$

Moreover, conditionally on \mathcal{T} , terminal node parameters are assumed independent: $p(\boldsymbol{\beta}|\mathcal{T}) = \prod_{r=1}^n p(\boldsymbol{\beta}_r|\mathcal{T})$. Therefore, the prior is fully determined by a tree prior $p(\mathcal{T})$, terminal node parameters priors $p(\boldsymbol{\beta}_r|\mathcal{T})$, and variance parameters priors $p(\sigma^2)$ and $p(\varphi_D)$.

3.1. *Tree prior.* The tree prior $p(\mathcal{T})$ is implicitly described by the stochastic tree-generating process of Chipman, George and McCulloch (1998), where each new tree is generated according to the following: (i) the probability for a node at depth q to be nonterminal, given by $\alpha(1 + q)^{-\nu}$, ($q = 1, 2, \dots$), (ii) the probability for a node to split at a predictor $x_{.j}$, ($j = 1, \dots, p$), given by the discrete uniform distribution on the set of available predictors, and (iii) given the predictor $x_{.j}$, the probability for a node to split at a value a , given by the discrete uniform distribution on the set of available splitting values. Probability (i) is a decreasing function of q , making deeper nodes less likely to split and favoring “bushy” trees. Chipman, George and McCulloch (1998) give guidelines to choose parameters α and ν by plotting the marginal prior distribution of the number of terminal nodes. In (ii) and (iii), predictors and splits are available if they lead to nonempty child nodes.

3.2. *Terminal node splines coefficients prior.* We follow Lang and Brezger (2004) and consider a conditionally conjugate P-spline prior: $\beta_r | \mathcal{T}, \tau^2 \propto \exp(-\frac{1}{2\tau^2} \beta_r' K_\beta \beta_r)$, where τ^2 is an additional smoothing variance parameter and

$$(3.2) \quad K_\beta = \begin{pmatrix} 1 & -1 & & & & \\ -1 & 2 & -1 & & & \\ & \ddots & \ddots & \ddots & & \\ & & -1 & 2 & -1 & \\ & & & -1 & 1 & \end{pmatrix}$$

is a penalty matrix of size $(m_D + \delta) \times (m_D + \delta)$, corresponding to a first order random walk. Note that this prior is improper, as the matrix K_β is not of full rank. In order to work with a proper prior in a model comparison setting, we replace the first and last element of the diagonal with $1 + \eta$, where η is a small constant. The model is completed by assigning a conjugate Inverse-Gamma hyperprior to the smoothing parameter $\tau^2 | \mathcal{T} \sim \text{IG}(a_\tau, b_\tau)$.

3.3. *Variance components priors.* We assume $\sigma^2 \sim \text{IG}(a_\sigma, b_\sigma)$. For φ_D , we choose the conjugate prior described by Rowe (2003) for autoregressive covariance matrices, with truncated support on $[0, 1]$. Let $0 = d_1 < \dots < d_{n_D} = D$ be the dose-escalation design sequence:

$$(3.3) \quad p(\varphi_D) \propto (1 - \varphi_D^2)^{-(n_D-1)/2} \exp\left(-\frac{\lambda_{01} - \varphi_D \lambda_{02} + \varphi_D^2 \lambda_{03}}{2(1 - \varphi_D^2)}\right) \mathbb{I}_{\varphi_D \in [0, 1]}$$

where \mathbb{I} is the indicator function, $\Lambda = (\Lambda_{vv'})_{1 \leq v, v' \leq n_D}$ is a hyperparameter matrix, and $(\lambda_{01}, \lambda_{02}, \lambda_{03})$ are defined through its diagonal, subdiagonal, and super-diagonal elements as follows: $\lambda_{01} = \sum_{v=1}^{n_D} \Lambda_{vv}$, $\lambda_{02} = \sum_{v=1}^{n_D-1} (\Lambda_{vv+1} + \Lambda_{v+1v})$, $\lambda_{03} = \sum_{d=2}^{n_D-1} \Lambda_{vv}$. In practice, we choose $\Lambda = Id_{n_D}$, the identity matrix of size $n_D \times n_D$, to put more weight on low values of φ_D and assume weak prior correlations between responses at different doses. This last distribution completes the prior model. We now turn to posterior inference on parameters, given the observations.

4. Posterior inference through MCMC simulation. We are interested in the posterior distribution $p(\mathcal{T}, \beta, \sigma^2, \varphi_D, \tau^2 | \mathbf{y})$. The rest of this section describes a Markov chain Monte Carlo algorithm for sampling from this distribution, as the number of potential trees prevents direct calculations. Our Gibbs sampler is adapted from the algorithms of Chipman, George and McCulloch (1998) and Gramacy and Lee (2008), with changes due to the specific structure of our model.

At each iteration, the algorithm performs a joint update of (\mathcal{T}, β) , conditionally on the rest of the parameters, followed by standard Gibbs component-wise updates

of each variance parameter. The joint tree and terminal nodes' spline coefficients update is decomposed into

$$(4.1) \quad \mathcal{T} | \mathbf{y}, \sigma^2, \varphi_D, \tau^2; \quad \text{followed by}$$

$$(4.2) \quad \beta_r | \mathcal{T}, \mathbf{y}, \sigma^2, \varphi_D, \tau^2; \quad \text{for } r \in \{1, \dots, n\}.$$

The draw of \mathcal{T} in (4.1) is performed by the Metropolis–Hastings algorithm of Chipman, George and McCulloch (1998), which simulates a Markov chain of trees that converges to the posterior distribution $p(\mathcal{T} | \mathbf{y}, \sigma^2, \varphi_D, \tau^2)$. The proposal density suggests a new tree based on four moves: grow a terminal node, prune a pair of terminal nodes, change the split rule of an internal node, and swap the splits of an internal node and one of its children's.

The target distribution can be decomposed as follows:

$$(4.3) \quad p(\mathcal{T} | \mathbf{y}, \sigma^2, \varphi_D, \tau^2) \propto p(\mathcal{T}) \int p(\mathbf{y} | \beta, \mathcal{T}, \sigma^2, \varphi_D, \tau^2) p(\beta | \mathcal{T}, \sigma^2, \varphi_D, \tau^2) d\beta.$$

The expression for the integral above is given in Low-Kam et al. (2015), in a closed form by conjugacy of the prior on $\beta = \{\beta_1, \dots, \beta_n\}$. Therefore, the draw of \mathcal{T} in (4.1) does not require a reversible-jump procedure for spaces of varying dimensions, even if nodes are added or deleted. The proposal density of the Metropolis–Hastings algorithm can be conveniently coupled with $p(\mathcal{T})$ to simplify calculations [Chipman, George and McCulloch (2002)]. Full conditional distributions for β_1, \dots, β_n in (4.2) and variance parameters σ^2, φ_D and τ^2 are available in Low-Kam et al. (2015).

Given posterior samples, predictive statistics are easily obtained via Monte Carlo simulation of $p(\mathbf{y}_i^* | \mathbf{y})$, for $i = 1, \dots, I$. More precisely, let $\mathbf{x}_i^* = \mathbf{x}_i$. At each iteration $\ell = 1, \dots, N$, the MCMC algorithm performs a draw from $p(\mathcal{T}, \beta, \sigma^2, \varphi_D, \tau^2 | \mathbf{y})$, followed by a draw of $\mathbf{y}_i^{(\ell)*}$ from the multivariate normal distribution $p(\mathbf{y}_i^{(\ell)*} | \mathcal{T}, \beta, \sigma^2, \varphi_D, \tau^2)$. In our case studies (Section 6), for example, we compare posterior summaries from the predictive distribution $p(y_{ik}^*(d) | \mathbf{y})$ to observed dose-response data $y_{ik}(d)$. We perform two series of posterior predictive checks: in the first one, the generated predictive samples are conditioned on the full set of dose-response curves, via the tree. The objective is to assess model adequacy and calibration. The second series studies model prediction accuracy using a leave-a-curve-out validation scheme, where each data curve is compared to the corresponding predictive sample obtained by fitting the tree on the remaining curves.

Posterior inference based on Monte Carlo samples is also used to derive inferential summaries about nontrivial functionals of the parameter/model space. The marginal effect of a physico-chemical property $x_{.j}$ on the response can be represented by the partial dependence function of Friedman (2001): let x_{1j}, \dots, x_{Sj} be a grid of new values for $x_{.j}$. Then the partial dependence function is $f(x_{sj}, d, t) =$

$(\sum_{i=1}^I f((x_{i1}, \dots, x_{ij-1}, x_{sj}, x_{ij+1}, \dots, x_{ip}), d, t))/I$, where x_{ij} is the i th observation of $x_{.j}$ in the data. For all doses, plotting the average of this function over Monte Carlo draws provides a visualization of the marginal effect of $x_{.j}$. This partial dependence function can also be extended to account for the joint marginal effect of two variables.

Similarly, posterior realizations $y|x$ can be used to report importance scores for each variable. For all $j \in \{1, \dots, P\}$, $S_j = \frac{\text{Var}\{\mathbb{E}\{y|x_{.j}\}\}}{\text{Var}\{y\}}$ and $T_j = \frac{\mathbb{E}\{\text{Var}\{y|x_{.-j}\}\}}{\text{Var}\{y\}}$ are the *first-order* and *total* sensitivity indices for variable $x_{.j}$, and represent the main and total influence, respectively, of this variable on the response [Gramacy, Taddy and Wild (2013)]. Unlike other metrics such as the variance reduction attributed to splits on the variable, sensitivity indices are robust to leaf model specifications and are therefore adapted for a dose-response leaf model. Both indices are defined given an uncertainty distribution on the inputs, usually the uniform distribution on the covariates space. We follow Gramacy, Taddy and Wild (2013) and use a Monte Carlo scheme to approximate S_j and T_j , that accounts for unknown responses by using predicted values for a Latin hypercube sampling design.

5. Extending the model to two-dimensional toxicity profiles. More general exposure escalation protocols involve the observation of a biological outcome y in association with a prescription of dose escalation $d \in [0, D]$, observed for a series of exposure times $t \in [t_0, T]$. Letting k , ($k = 1, \dots, K$) be a replication index, we define $y_{ik}(d, t)$ as the outcome of interest, evaluated at dose d , time t and extend the model in (2.1): $y_{ik}(d, t) = f(\mathbf{x}_i, d, t) + \varepsilon_{ik}(d, t)$, where f is a random mean response surface and $\varepsilon_{ik}(d, t) \sim N(0, \sigma_{dt}^2)$. To account for dependence between doses and durations of exposure, for each nanoparticle i , we assume $\text{Cov}(\varepsilon_{ik}(d, t), \varepsilon_{ik'}(d', t')) = \sigma^2 \varphi_D^{|d-d'|} \varphi_T^{|t-t'|}$, where $\varphi_D \in [0, 1]$ and $\varphi_T \in [0, 1]$ are autocorrelation parameters.

The response surface f in the terminal nodes of \mathcal{T} is modeled by a tensor product of two one-dimensional P-splines [Lang and Brezger (2004)]. Let $\mathcal{B}_1(\cdot), \dots, \mathcal{B}_{m_D+\delta}(\cdot)$ defined as in Section 2.2 and $\mathcal{B}_1(\cdot), \dots, \mathcal{B}_{m_T+\zeta}(\cdot)$ denote $m_T + \zeta$ B-spline basis functions of order ζ on $[t_0, T]$, with m_T fixed knots. Then, if \mathbf{x}_i is in the subset corresponding to the r th terminal node of T_j , $f(\mathbf{x}_i, d, t) = \sum_{\ell=1}^{m_D+\delta} \sum_{m=1}^{m_T+\zeta} \beta_{r\ell m} \mathcal{B}_\ell(d) \mathcal{B}_m(t)$, where $\beta_r = (\beta_{r11}, \dots, \beta_{r(m_D+\delta)(m_T+\zeta)})'$ is a vector of spline coefficients associated to the r th terminal node.

The prior model has the same global dependence structure as in Section 3, but now includes an additional independent term φ_T for time-covariance. Let $t_0 = t_1 < \dots < t_{n_T} = T$ be the sequence of exposure times when toxicity was measured. We adapt prior (3.3) to preserve conjugacy and introduce a similar distribution for φ_T :

$$(5.1) \quad p(\varphi_D) \propto (1 - \varphi_D^2)^{-(n_T(n_D-1))/2} \exp\left(-\frac{\lambda_{01} - \varphi_D \lambda_{02} + \varphi_D^2 \lambda_{03}}{2(1 - \varphi_D^2)}\right) \mathbb{I}_{\varphi_D \in [0, 1]},$$

$$(5.2) \quad p(\varphi_T) \propto (1 - \varphi_T^2)^{-(n_D(n_T-1))/2} \exp\left(-\frac{\gamma_{01} - \varphi_T \gamma_{02} + \varphi_T^2 \gamma_{03}}{2(1 - \varphi_T^2)}\right) \mathbb{I}_{\varphi_T \in [0, 1]},$$

where γ_{01} , γ_{02} and γ_{03} are obtained by summing elements of the diagonal, sub-diagonal, and superdiagonal of matrix parameter prior Γ , constructed following the guidelines introduced in Section 3.3. For the terminal nodes' spline coefficient priors, we use a spatial extension of [Besag and Kooperberg \(1995\)](#), a first order random walk prior based on the four nearest neighbours of splines coefficients, with appropriate changes for corners and edges: $\beta_r | \mathcal{T}, \tau^2 \propto \exp(-\frac{1}{2\tau^2} \beta_r' K_\beta \beta_r)$, where K_β is a penalty band matrix of size $(m_D + \delta)(m_T + \zeta) \times (m_D + \delta)(m_T + \zeta)$, which extends matrix (3.2) to the two-dimensional case. For posterior inference, we add a step to generate φ_T in the Gibbs sampler of Section 4.

6. Applications. A simulation study to assess model performance is described in [Low-Kam et al. \(2015\)](#). In the rest of this section we illustrate our approach with experimental results from a case study reported by [Zhang et al. \(2012\)](#), measuring toxicity of 24 metal oxides on human bronchial epithelial (BEAS-2B) cells.

6.1. *Case studies background.* After 24 h, Lactate Dehydrogenase (LDH) release was used to measure the death rate of cells exposed to eleven doses of metal oxides (from 0 to 200 $\mu\text{g/ml}$), evenly spaced on the logarithmic scale. Cell death is commonly used to screen for ENM cytotoxicity without reference to a specific mechanism. Figure 1 shows the LDH dose-responses curves for the 24 metal oxide nanoparticles. In a second assay, Propidium Iodide (PI) fluorescence was used to indicate the percentage of cells experiencing oxidative stress through cellular surface membrane permeability, across the same ten doses and after six times of exposure (from 1 to 6 h, at every hour). Figure 2 shows a heatmap representation for the PI assay, for all metal oxides, doses, times, and replicates, where responses are color-coded from light (low) to dark (high). In both assays, seven metal oxides (Co_3O_4 , CoO , Cr_2O_3 , CuO , Mn_2O_3 , Ni_2O_3 and ZnO) display a notable rise for the higher doses, suggesting toxicity.

All metal oxides are characterized by six physico-chemical properties of potential interest to explain toxicity profiles: nanoparticle size in media, a measure of the crystalline structure ($b(\text{\AA})$), lattice energy ($\Delta H_{\text{lattice}}$), which measures the strength of the bonds in the nanoparticles, the enthalpy of formation ($\Delta H_{Me^{n+}}$), which is a combined measure of the energy required to convert a solid to a gas and the energy required to remove n electrons from that gas, metal dissolution rate, and conduction band energy (the energy to free electrons from binding with atoms).

In our analysis, we use cubic splines, that is, $\delta = \zeta = 4$, and place interior knots at each intermediate dose from 0.39 to 100 $\mu\text{g/ml}$. Therefore, $n_D = m_D + 2$ and $n_T = m_T + 2$. For the tree prior, we adopt the default choice of [Chipman, George and McCulloch \(1998\)](#), $(\alpha, \nu) = (0.95, 2)$, which puts more weight on trees of size 2 or 3. We place relatively diffuse priors $\text{Gamma}(1, 1)$ on precision parameters $1/\tau^2$ and $1/\sigma^2$. We choose $\Lambda = Id_{n_D}$ and $\Gamma = Id_{n_T}$, assuming no prior correlations between measurements at different doses and times. Finally, moves ‘‘Grow,’’

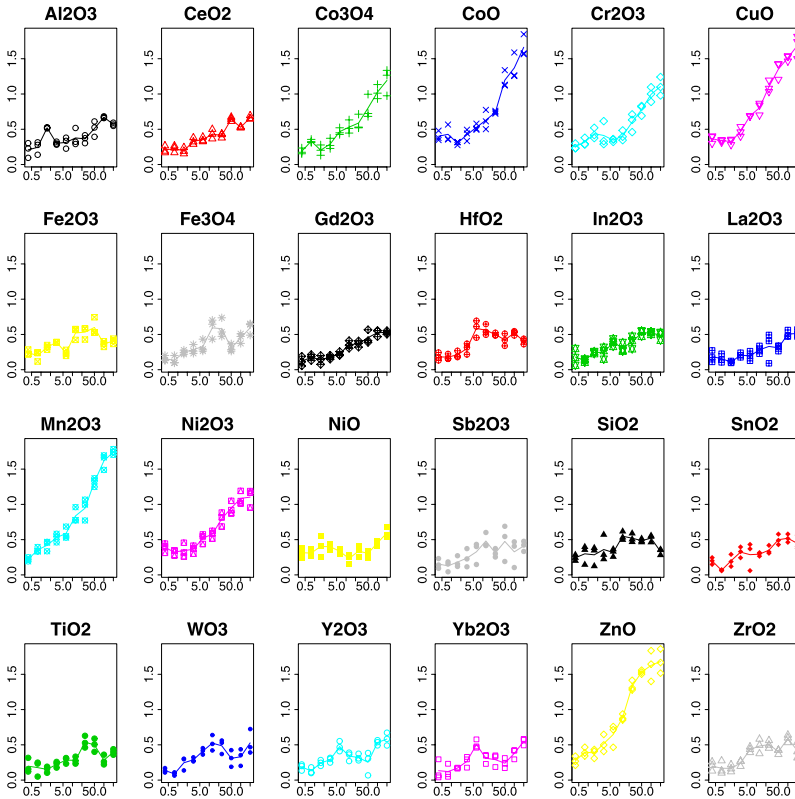


FIG. 1. Dose-response curves for LDH assay.

“Prune,” “Change” and “Swap” of the Metropolis–Hastings tree-generating algorithm have probabilities 0.1, 0.1, 0.6 and 0.2, respectively. We used a total of 160,000 iterations. After discarding 80,000 iterations for burn-in, the remaining samples for estimation were thinned to save computer storage. The rest of this section shows the results obtained on LDH and PI assays.

6.2. LDH dose-escalation assay. Figure 3 (top) shows both sensitivity indices described in Section 4 for the six physico-chemical properties. Figure 3 (bottom) shows the combined marginal effect of conduction band energy and dissolution on LDH, obtained with the partial dependence function of Friedman (2001), and color-coded from light (low) to dark (high), for dose 200 $\mu\text{g/ml}$. The tree isolates a first region of high toxicity, corresponding to ENM with high dissolution rates (ZnO and CuO). This region corresponds to the first mechanism of toxicity identified by Zhang et al. (2012): highly soluble metal oxides, such as ZnO and CuO, are more likely to release metal ions and disturb the cellular state. A second region of toxicity on Figure 3 (left) includes metal oxides Co_3O_4 , CoO, Cr_2O_3 , Mn_2O_3 and Ni_2O_3 , with E_c values ranging from -4.33 eV for Mn_2O_3 to -4.59 eV for Ni_2O_3 .




FIG. 2. Heatmap for PI fluorescence assay, color-coded from light (low) to dark (high). Each row corresponds to a nanoparticle at one dose across 6 times (1 to 6 h) and 4 replicates. For each nanoparticle, there are 11 rows, one for each dose (0 to 200 µg/ml), arranged from bottom to top.

This region matches the second mechanism for toxicity described by Zhang et al. (2012): the overlap of the conduction band energy of the metal oxides with the biological redox potential of cells, ranging from -4.12 to -4.84 eV. When these two energy levels are alike, transfer of electrons from metal oxides to cells is facilitated, disturbing the intracellular state. Note that Figure 3 (bottom) also shows an additional split that isolates Mn₂O₃, whose toxicity for the LDH assay is more comparable to ZnO and CuO (see Figure 1). Similar figures for other doses are included in Low-Kam et al. (2015). The LDH assay illustrates how threshold effects and interactions of physico-chemical properties are accurately captured by a tree structure.

We perform posterior predictive checks for model fitting. Figure 4 shows the expected posterior predictive dose-response curves for two nontoxic metal oxides (CeO₂ and Fe₃O₄) and two toxic ones (Cr₂O₃ and ZnO), with the associated 90% intervals. All four intervals provide good coverage for the original data. The other 20 curves exhibit similar behavior and can be found in Low-Kam et al. (2015). We also study the prediction accuracy of the model using a leave-a-curve-out validation framework. Results for CeO₂, Fe₃O₄, Cr₂O₃ and ZnO are presented in Low-Kam et al. (2015). While leave-one-out predictions recover general trends, in some cases we observe suboptimal coverage, especially in sparse areas of the physico-chemical spectrum. For example, nanoparticles ZnO and CuO alone determine tree splits on the metal dissolution parameter and, once removed, cannot be accurately predicted by the model.

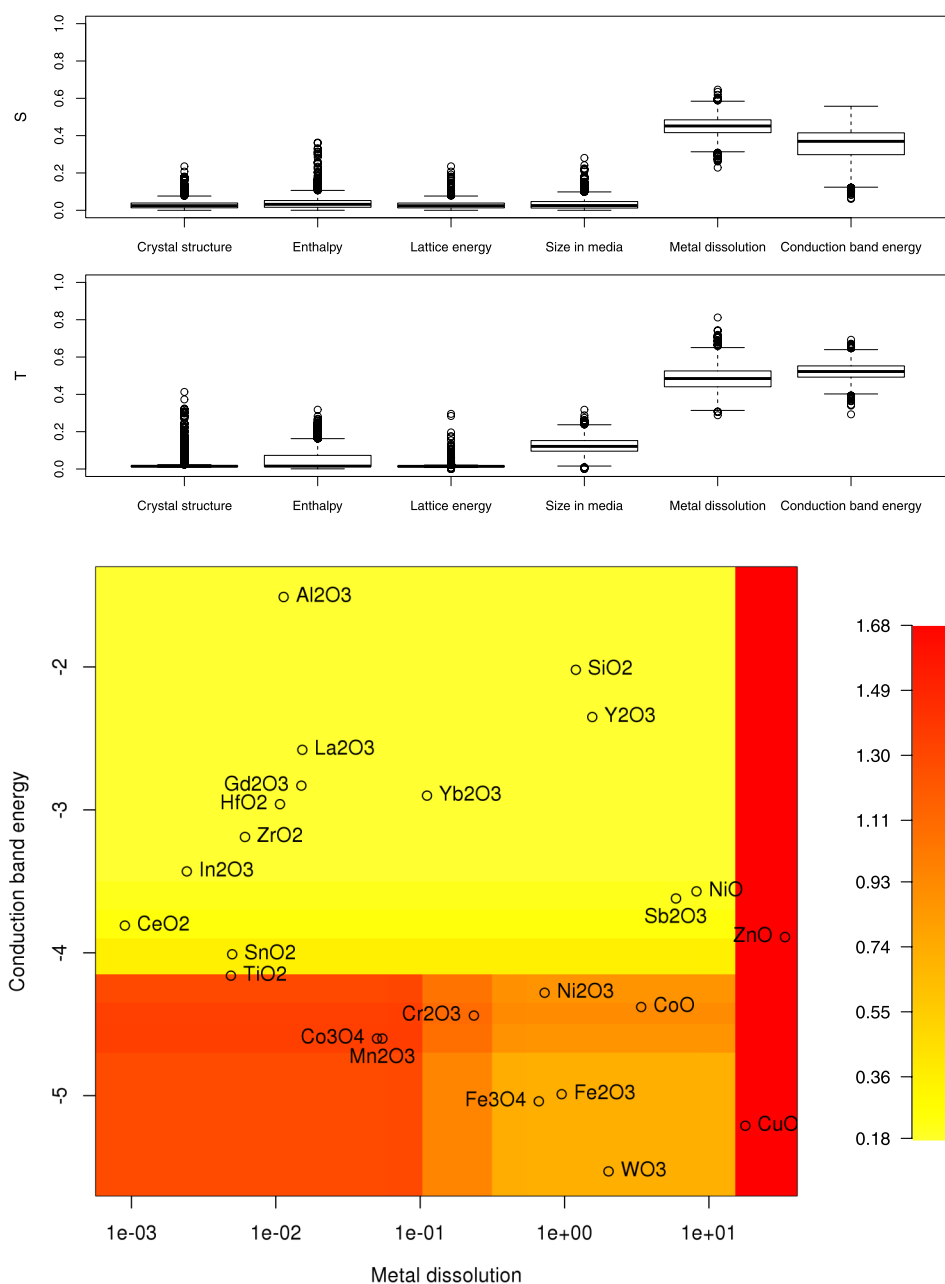


FIG. 3. LDH assay. (Top) First order (S) and total (T) sensitivity indices for the six physio-chemical properties in the LDH assay. (Bottom) 2-dimensional partial dependence function for marginal effect of metal dissolution (log scale) and conduction band energy in the LDH assay at 200 $\mu\text{g}/\text{ml}$. The toxicity response is color-coded from light (low) to dark (high). The figure also shows the projections of the 24 metaloxides in this subspace.

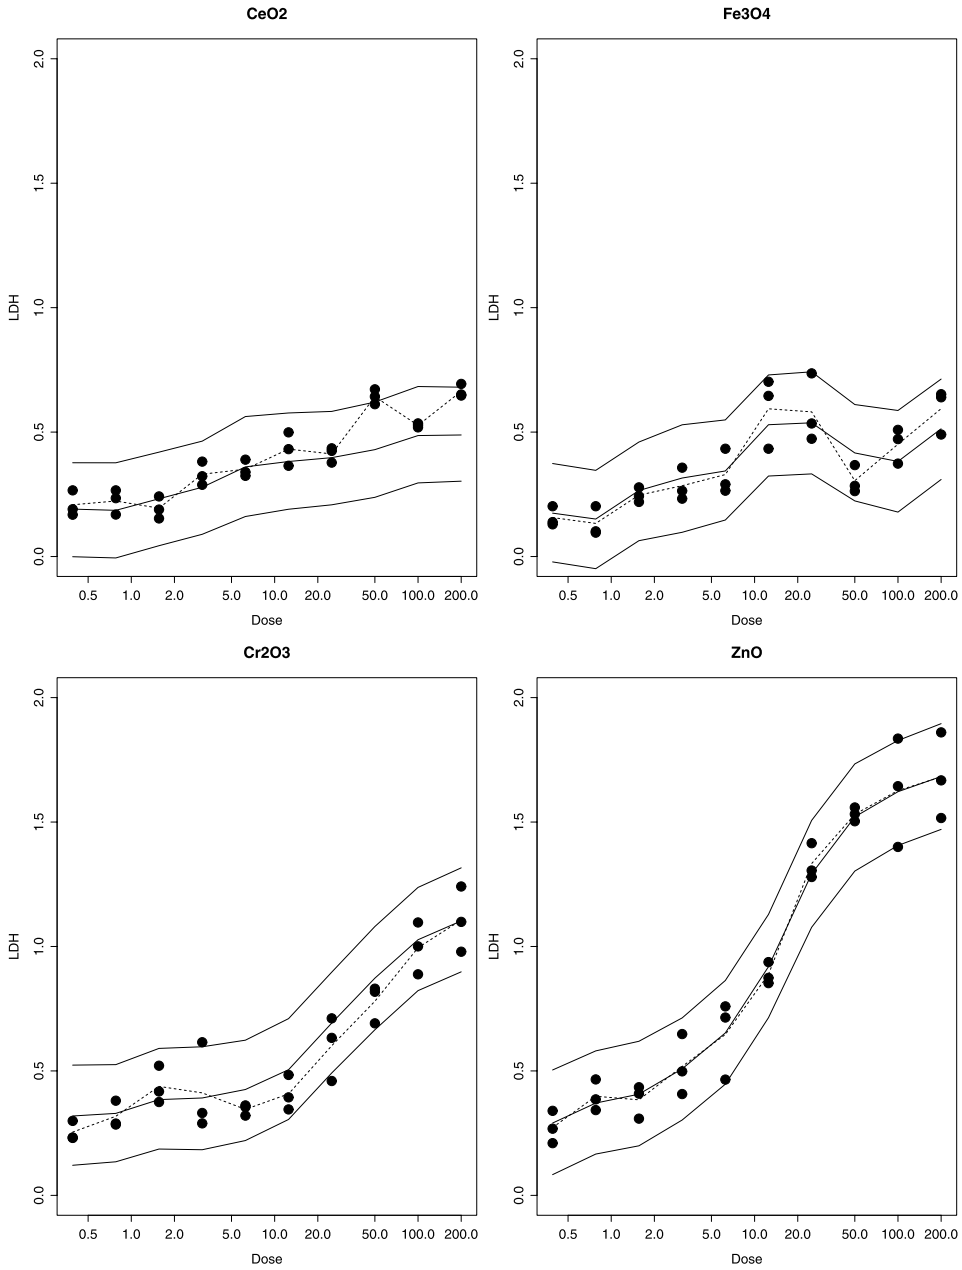


FIG. 4. Posterior predictive curves for CeO_2 , Fe_3O_4 , Cr_2O_3 and ZnO . The points are the observed replicates and the dashed line is the average observed response. The expected posterior predictive curve and 90% interval are in solid lines.

Finally, the proposed methodology is compared for validation to the Bayesian Additive Regression Trees (BART) method of Chipman, George and McCulloch (2010a), a sum-of-tree extension of Chipman, George and McCulloch (1998), with the R package “BayesTree” [Chipman, George and McCulloch (2010b)]. As BART model one-dimensional responses, we use the area under the LDH curves (AUC) as the dependent variable. In Chipman, George and McCulloch (2010a), the proportion of all splitting rules attributed to a variable at each draw on all trees, averaged over all iterations, is proposed as a measure of variable importance, when the number of trees is small. Results are presented in Low-Kam et al. (2015). Variable importance scores and marginal effects from BART are similar to those obtained with our method and confirm that the AUC is an accurate summary for toxicity for the LDH assay. The first advantage of using a dose-response leaf model instead of the AUC is that we avoid preliminary assessment of the data for choosing a summary over another: toxicologists usually report several toxicity parameters (EC50, slope), as they may convey different information. The second advantage is better understood from a predictive perspective, as our model allows for full dose-response dynamics instead of the AUC. A comparison with the Treed Gaussian Process, using the R package `tgp` [Gramacy and Taddy (2010)], is also included in Low-Kam et al. (2015). After tuning `tgp` to forbid splitting on dose (`basemax`, `splitmin`), we can indeed reproduce the essential structure of our model using this well-tested R library. Our findings proved to be robust to differing details in the prior specification, as the model fit with `tgp` also captures the marginal effects of the predictors metal dissolution and conduction band energy on toxicity.

6.3. *PI general exposure assay.* Figure 5 (top) shows the variable sensitivity indices of the six physico-chemical properties. Figure 5 (bottom) illustrates the marginal effect of both conduction band energy and dissolution on membrane damage, calculated with the partial dependence function, and color-coded from light to dark, for dose 200 $\mu\text{g}/\text{ml}$ and time 6 h. The tree model for PI assay also identifies the two areas of toxicity indicated in Zhang et al. (2012), corresponding to highly soluble metal oxides and nanoparticles whose conduction band energy overlaps with cellular redox potential range. Additional figures for marginal effect of conduction band energy and metal dissolution, for all doses and times, are included in Low-Kam et al. (2015). The similarity of variable importance scores and marginal effect of conduction band energy and dissolution obtained for LDH and PI assays indicates a strong correlation between these assays for nanoparticle toxicity assessment, as noted by Zhang et al. (2012). Figure 6 illustrates the posterior predictive 90% surface intervals for two nontoxic metal oxides (La_2O_3 and TiO_2) and two toxic ones (Co_3O_4 and CuO), showing good posterior coverage over all doses and times of exposure. Similar surfaces for the other 20 metal oxides are plotted in Low-Kam et al. (2015). Leave-a-surface-out predictions for La_2O_3 , TiO_2 , Co_3O_4 , and CuO are presented in the appendix and show the limitations of the model for prediction when extrapolating to sparse areas of the covariate space, similar to what we observed in the LDH assay.

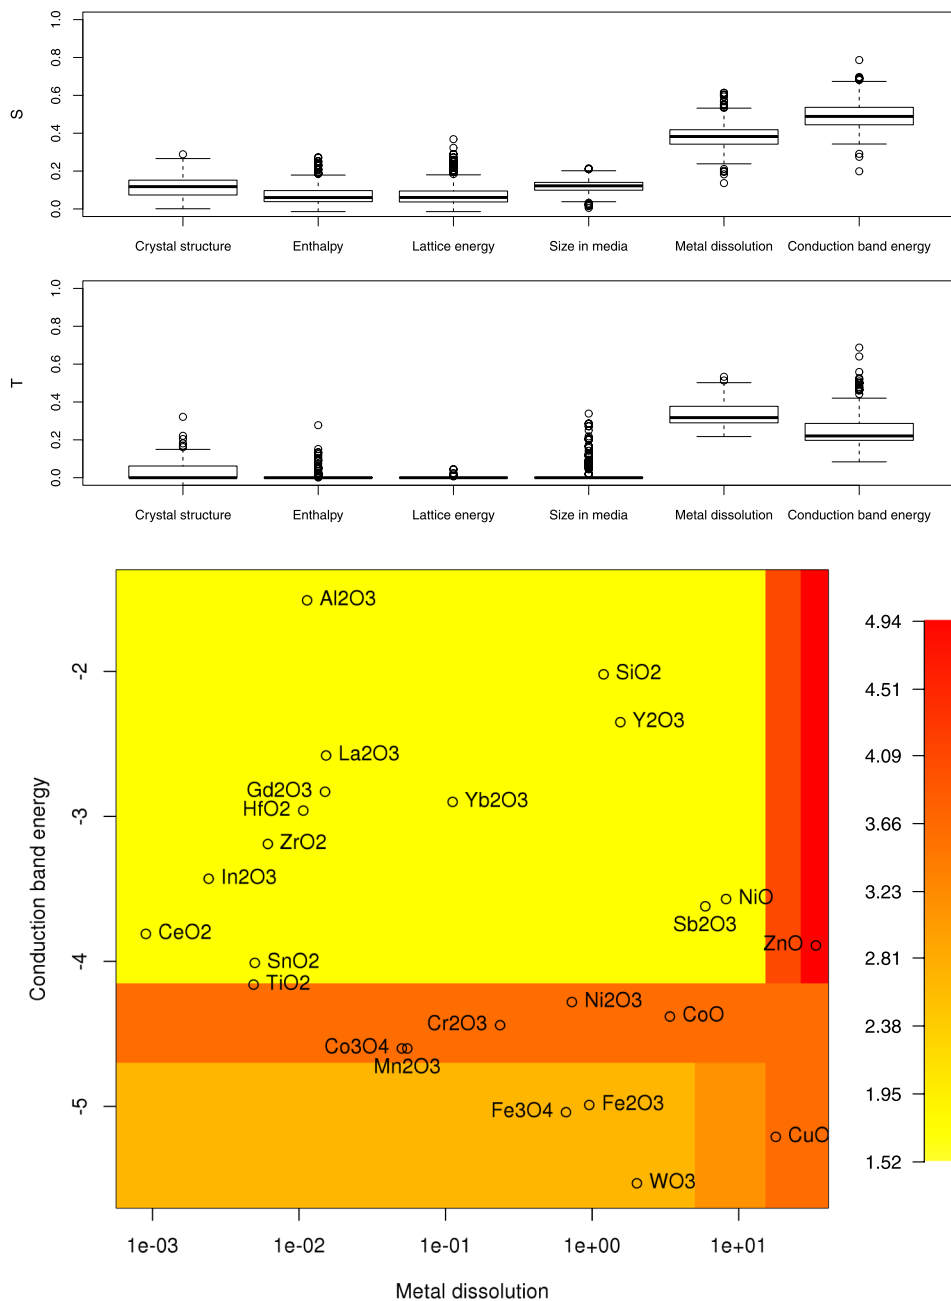


FIG. 5. PI uptake. (Top) First order (*S*) and total (*T*) sensitivity indices for the six physio-chemical properties in the PI assay. (Bottom) 2-dimensional partial dependence function for marginal effect of metal dissolution (log scale) and conduction band energy in the PI assay at 200 $\mu\text{g}/\text{ml}$ and 6 h. The toxicity response is color-coded from light (low) to dark (high). The figure also shows the projections of the 24 metaloxides in this subspace.

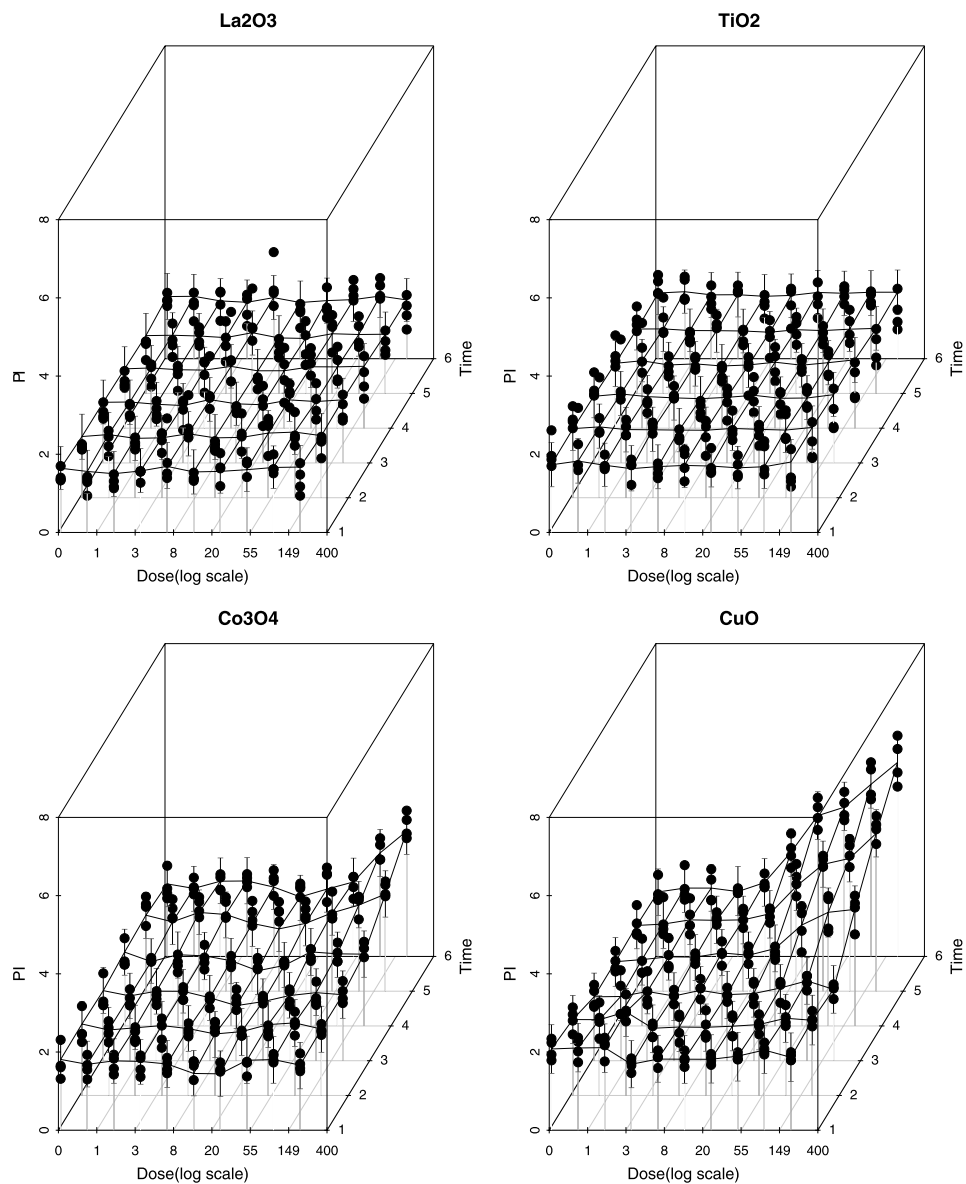


FIG. 6. Posterior predictive surfaces for La_2O_3 , TiO_2 , Co_3O_4 and CuO . The solid line is the expected posterior predictive surface with the associated 90% interval. The points are the observed data replicates.

7. Discussion. We propose a Bayesian regression tree model to define relationships between physico-chemical properties of engineered nanomaterials and their functional toxicity profiles in dose-escalation assays. As demonstrated by the case studies, the tree structure is adapted to account for flexible models of

structure-activity relationships, such as threshold effects and interactions. The proposed model integrates information across all doses and replicates, and therefore is adapted to small sample sizes usually found in nanotoxicology data sets. Monte Carlo integration over the model space provides straightforward inference on non-trivial functionals of parameters of interest and prediction of full dose-response curves from nanoparticle characteristics. The smoothing splines representation allows for easy extension of the model to two-dimensional toxicity profiles of general exposure escalation assays as well as for modeling sub-lethal outcomes.

The convergence of Bayesian tree models should be carefully assessed for all applications of the proposed methodology. The four moves of the Metropolis–Hastings algorithm of Chipman, George and McCulloch (1998) work well in our simulations and case studies, however, other applications might require additional moves to move faster through the tree space and improve convergence [see, e.g., Gramacy and Lee (2008), Wu, Tjelmeland and West (2007)]. As illustrated in Section 6, another potential pitfall of the model is its predictive performance for sparsely explored nanoparticle characteristics. This issue is not specific to our model and possible improvements would be obtained by combining multiple studies in a meta-analysis framework, with the appropriate adjustments for data heterogeneity or formalizing explicit prior knowledge about hazardous nanoparticle properties.

As seen in the case study for cell death and cellular membrane permeability, different toxicity mechanisms can be closely related. Therefore, an important opportunity for model extensions would be to combine different biological assays in a single analysis, the final goal being that of understanding if nanoparticles physical and chemical properties have a differential effect on different cellular injury pathways. This would require more sophisticated modeling strategies that will be more likely to be useful if technological advances will allow for feasible screening of much larger nanomaterial libraries.

Acknowledgments. Any opinions, findings, conclusions or recommendations expressed herein are those of the author(s) and do not necessarily reflect the views of the National Science Foundation or the Environmental Protection Agency. This work has not been subjected to an EPA peer and policy review.

SUPPLEMENTARY MATERIAL

Additional results for online publication (DOI: [10.1214/14-AOAS797SUPPA](https://doi.org/10.1214/14-AOAS797SUPPA); .pdf). This appendix provides full conditional distributions and additional experimental results.

Code (DOI: [10.1214/14-AOAS797SUPPB](https://doi.org/10.1214/14-AOAS797SUPPB); .zip). This folder contains a C++ implementation of the algorithm.

REFERENCES

- BESAG, J. and KOOPERBERG, C. (1995). On conditional and intrinsic autoregressions. *Biometrika* **82** 733–746. [MR1380811](#)
- BREIMAN, L., FRIEDMAN, J. H., OLSHEN, R. A. and STONE, C. J. (1984). *Classification and Regression Trees*. Wadsworth, Belmont, CA. [MR0726392](#)
- CHIPMAN, H. A., GEORGE, E. I. and MCCULLOCH, R. E. (1998). Bayesian CART model search. *J. Amer. Statist. Assoc.* **93** 935–948.
- CHIPMAN, H. A., GEORGE, E. I. and MCCULLOCH, R. E. (2002). Bayesian treed models. *Machine Learning* **48** 299–320.
- CHIPMAN, H. A., GEORGE, E. I. and MCCULLOCH, R. E. (2010a). BART: Bayesian additive regression trees. *Ann. Appl. Stat.* **4** 266–298. [MR2758172](#)
- CHIPMAN, H. A., GEORGE, E. I. and MCCULLOCH, R. E. (2010b). Implementation of BART: Bayesian additive regression trees. R package version 0.3-1.1.
- DE'ATH, G. (2002). Multivariate regression trees: A new technique for modeling species-environment relationships. *Ecology* **83** 1105–1117.
- EILERS, P. H. C. and MARX, B. D. (1996). Flexible smoothing with B-splines and penalties. *Statist. Sci.* **11** 89–121. [MR1435485](#)
- FRIEDMAN, J. H. (2001). Greedy function approximation: A gradient boosting machine. *Ann. Statist.* **29** 1189–1232. [MR1873328](#)
- GALIMBERTI, G. and MONTANARI, A. (2002). Regression trees for longitudinal data with time-dependent covariates. In *Classification, Clustering, and Data Analysis* 391–398. Springer, Berlin. [MR2010476](#)
- GRAMACY, R. B. and LEE, H. K. H. (2008). Bayesian treed Gaussian process models with an application to computer modeling. *J. Amer. Statist. Assoc.* **103** 1119–1130. [MR2528830](#)
- GRAMACY, R. B. and TADDY, M. A. (2010). Categorical inputs, sensitivity analysis, optimization and importance tempering with tgp version 2, an R package for treed Gaussian process models. *Journal of Statistical Software* **33** 1–48.
- GRAMACY, R. B., TADDY, M. and WILD, S. M. (2013). Variable selection and sensitivity analysis using dynamic trees, with an application to computer code performance tuning. *Ann. Appl. Stat.* **7** 51–80. [MR3086410](#)
- KONOMI, B., KARAGIANNIS, G., SARKAR, A., SUN, X. and LIN, G. (2014). Bayesian treed multivariate Gaussian process with adaptive design: Application to a carbon capture unit. *Technometrics* **56** 145–158. [MR3207843](#)
- LANG, S. and BREZGER, A. (2004). Bayesian P-splines. *J. Comput. Graph. Statist.* **13** 183–212. [MR2044877](#)
- LIU, R., RALLO, R., GEORGE, S., JI, Z., NAIR, S., NEL, A. E. and COHEN, Y. (2011). Classification NanoSAR development for cytotoxicity of metal oxide nanoparticles. *Small* **7** 1118–1126.
- LOW-KAM, C., TELESKA, D., JI, Z., ZHANG, H., XIA, T., ZINK, J. I. and NEL, A. (2015). Supplement to “A Bayesian regression tree approach to identify the effect of nanoparticles’ properties on toxicity profiles.” DOI:10.1214/14-AOAS797SUPPA, DOI:10.1214/14-AOAS797SUPPB.
- PATEL, T., TELESKA, D., LOW-KAM, C., JI, Z. X., ZHANG, H. Y., XIA, T., ZINC, J. I. and NEL, A. E. (2014). Relating nano-particle properties to biological outcomes in exposure escalation experiments. *Environmetrics* **25** 57–68. [MR3233744](#)
- RAMSAY, J. O. (1998). Monotone regression splines in action. *Statist. Sci.* **3** 425–441.
- ROWE, D. B. (2003). *Multivariate Bayesian Statistics: Models for Source Separation and Signal Unmixing*. Chapman & Hall/CRC, Boca Raton, FL. [MR2000406](#)
- SEGAL, M. R. (1992). Tree-structured methods for longitudinal data. *J. Amer. Statist. Assoc.* **87** 407–418.
- SELA, R. J. and SIMONOFF, J. S. (2012). RE-EM trees: A data mining approach for longitudinal and clustered data. *Mach. Learn.* **86** 169–207. [MR2892116](#)

- WU, Y., TJELMELAND, H. and WEST, M. (2007). Bayesian CART: Prior specification and posterior simulation. *J. Comput. Graph. Statist.* **16** 44–66. [MR2345747](#)
- YU, Y. and LAMBERT, D. (1999). Fitting trees to functional data, with an application to time-of-day patterns. *J. Comput. Graph. Statist.* **8** 749–762.
- YU, K., WHEELER, W., LI, Q., BERGEN, A. W., CAPORASO, N., CHATTERJEE, N. and CHEN, J. (2010). A partially linear tree-based regression model for multivariate outcomes. *Biometrics* **66** 89–96. [MR2756694](#)
- ZHANG, S., SHIH, Y.-C. T. and MÜLLER, P. (2007). A spatially-adjusted Bayesian additive regression tree model to merge two datasets. *Bayesian Anal.* **2** 611–633. [MR2342177](#)
- ZHANG, H., JI, Z., XIA, T., MENG, H., LOW-KAM, C., LIU, R., POKHREL, S., LIN, S., WANG, X., LIAO, Y.-P., WANG, M., LI, L., RALLO, R., DAMOISEAUX, R., TELESKA, D., MÄDLER, L., COHEN, Y., ZINK, J. I. and NEL, A. E. (2012). Use of metal oxide nanoparticle band gap to develop a predictive paradigm for oxidative stress and acute pulmonary inflammation. *ACS Nano* **6** 4349–4368.

C. LOW-KAM
D. TELESKA
DEPARTMENT OF BIOSTATISTICS
UNIVERSITY OF CALIFORNIA
LOS ANGELES, CALIFORNIA 90095
USA
E-MAIL: clowkam@gmail.com
dtelesca@ucla.edu

T. XIA
DIVISION OF NANOMEDICINE
DEPARTMENT OF MEDICINE
UNIVERSITY OF CALIFORNIA
LOS ANGELES, CALIFORNIA 90095
USA
E-MAIL: txia@ucla.edu

Z. JI
H. ZHANG
CALIFORNIA NANOSYSTEMS INSTITUTE
UNIVERSITY OF CALIFORNIA
LOS ANGELES, CALIFORNIA 90095
USA
E-MAIL: zji@cnsi.ucla.edu
zhangh@ucla.edu

J. I. ZINK
DEPARTMENT OF CHEMISTRY & BIOCHEMISTRY
UNIVERSITY OF CALIFORNIA
LOS ANGELES, CALIFORNIA 90095
USA
E-MAIL: zink@chem.ucla.edu

A. E. NEL
CALIFORNIA NANOSYSTEMS INSTITUTE
UNIVERSITY OF CALIFORNIA
LOS ANGELES, CALIFORNIA 90095
USA
AND
DIVISION OF NANOMEDICINE
DEPARTMENT OF MEDICINE
UNIVERSITY OF CALIFORNIA
LOS ANGELES, CALIFORNIA 90095
USA
E-MAIL: ANel@mednet.ucla.edu