



Published in final edited form as:

*Ann Appl Stat.* 2021 September ; 15(3): 1405–1430. doi:10.1214/21-aoas1461.

## BAYESIAN JOINT MODELING OF CHEMICAL STRUCTURE AND DOSE RESPONSE CURVES

Kelly R. Moran<sup>1,\*</sup>, David Dunson<sup>1,†</sup>, Matthew W. Wheeler<sup>2</sup>, Amy H. Herring<sup>1,‡</sup>

<sup>1</sup>Department of Statistical Science, Duke University

<sup>2</sup>Biostatistics and Computational Biology Branch, National Institute of Environmental Health Sciences

### Abstract

Today there are approximately 85,000 chemicals regulated under the Toxic Substances Control Act, with around 2,000 new chemicals introduced each year. It is impossible to screen all of these chemicals for potential toxic effects, either via full organism *in vivo* studies or *in vitro* high-throughput screening (HTS) programs. Toxicologists face the challenge of choosing which chemicals to screen, and predicting the toxicity of as yet unscreened chemicals. Our goal is to describe how variation in chemical structure relates to variation in toxicological response to enable *in silico* toxicity characterization designed to meet both of these challenges. With our Bayesian partially Supervised Sparse and Smooth Factor Analysis (BS<sup>3</sup>FA) model, we learn a distance between chemicals targeted to toxicity, rather than one based on molecular structure alone. Our model also enables the prediction of chemical dose-response profiles based on chemical structure (i.e., without *in vivo* or *in vitro* testing) by taking advantage of a large database of chemicals that have already been tested for toxicity in HTS programs. We show superior simulation performance in distance learning and modest to large gains in predictive ability compared to existing methods. Results from the high-throughput screening data application elucidate the relationship between chemical structure and a toxicity-relevant high-throughput assay. An R package for BS<sup>3</sup>FA is available online at <https://github.com/kelrenmor/bs3fa>.

### Key words and phrases

Dimension reduction; distance learning; functional prediction; high-throughput screening; toxicity; ToxCast; QSAR

---

\* kelly.r.moran@duke.edu . † dunson@duke.edu . ‡ amy.herring@duke.edu .

#### SUPPLEMENTARY MATERIAL

**Supplement to “Bayesian joint modeling of chemical structure and dose-response curves”** (DOI: [10.1214/21-AOAS1461SUPPA](https://doi.org/10.1214/21-AOAS1461SUPPA); .pdf). Additional explanatory text, sampler details, information on simulated and non-synthetic data runs, and results.

**Code** (DOI: [10.1214/21-AOAS1461SUPPB](https://doi.org/10.1214/21-AOAS1461SUPPB); .zip). Code used to clean data, run simulation, and perform ToxCast analysis in the manuscript.

**Data** (DOI: [10.1214/21-AOAS1461SUPPC](https://doi.org/10.1214/21-AOAS1461SUPPC); .zip). Data used to conduct ToxCast analyses in the manuscript.

## 1. Introduction.

Daily life involves being exposed to a variety of chemical substances from diverse sources and at varying concentrations. A myriad of legislation and regulatory bodies work to assess consumer and industrial products for toxicity and reduce exposure risk. The Toxic Substances Control Act (TSCA), passed by Congress in 1976 and administered by the U.S. Environmental Protection Agency (EPA), regulates the bulk of<sup>1</sup> new and existing chemicals in the United States (U.S.). When the TSCA was enacted, around 60,000 chemicals were grandfathered into the program and effectively considered safe for use. The EPA has struggled to catch up on this backlog while also keeping up with the rate of new introductions (roughly 2,000 chemicals per year), as they assess chemicals for potential toxicity. High-throughput screening methods have proved vital to this effort, as they allow researchers to quickly conduct millions of tests.

The EPA's Toxicity Forecaster (ToxCast) research program in which thousands of chemicals are tested in more than 700 high-throughput assay endpoints is used to prioritize, screen and evaluate chemicals for potential toxic effects (Dix et al. (2007), Judson et al. (2009), Kavlock et al. (2012)). However, even high-throughput toxicity screening (HTS) programs, which allow for the relatively cheap and fast collection of dose-response information via *in vitro* studies rather than full organism *in vivo* studies, are still too slow and expensive to be able to study all chemicals. *In silico* studies, that is, those performed via computer modeling rather than in the lab, can be used to guide the design of and supplement the results from lab-based studies. Specifically, the characterization of an activity relevant chemical distance *in silico* enables more targeted design of further *in vitro* studies, increasing the efficiency of resource allocation. In addition, predicting toxicity via such studies helps bridge the gap between the number of chemicals of interest and the number with known toxicological profiles.

The goal of this work is to make inferences about how variation in chemical structure relates to variation in toxicological response. Sparse function-on-scalars regression models (Barber, Reimherr and Schill (2017), Chen, Goldsmith and Ogden (2016), Fan and Reimherr (2017), Kowal and Bourgeois (2020)) do this in a limited way by selecting the important chemical structure features and giving them appropriate coefficients or weights. Because there are many redundant and highly correlated structure features (see Figure 1), a PCA-esque approach that introduces latent factors related to the major directions of variation in the molecular structure is more informative than such penalized regression approaches. However, simply performing PCA or other unsupervised dimension reduction approaches on the chemical structure ignores the distinction between *overall* variation and *toxicity-relevant* variation in the molecular structure. A supervised dimension reduction approach, on the other hand, provides a coherent and flexible framework within which to describe the relationship between molecular variation and activity variation via a shared latent subspace.

---

<sup>1</sup>Exceptions regulated under different legislation include foods and food additives, drugs, cosmetics, pesticides, tobacco products, research substances used in small quantities and radioactive materials and waste.

Dose response data and molecular structure are the two sources of information in ToxCast relevant to addressing our goal. Explicitly, chemical  $i$  in ToxCast has two relevant pieces of information: the vector of response observations at  $D$  doses  $y_i = [y_i(d_1), \dots, y_i(d_D)]'$ , and the vector of  $S$  molecular features  $x_i = [x_{i1}, \dots, x_{iS}]'$ . Observations  $y_i$  are sparse, noisy and not on a regular grid. For an example, see Figure 2; Chlorobenzilate has 54 observations at 11 unique doses, yet 5-Methyl-1H-benzotriazole only has three doses with one observation each. Not all aspects of the feature space (i.e., not all entries in  $x_i$ ) are likely to be relevant to the toxicological response.

Response variables in ToxCast are assay specific. Assays in ToxCast each measure a single endpoint, for example, the binding to a certain receptor protein, or the transcription of a target gene. The specific assay endpoint considered in our real-data example, the AttaGene pregnane X receptor (PXR) assay, records the fold-change values in the activity of the nuclear pregnane X receptor for drug-treated vs. control-treated human hepatic cells (specifically, a HG19 subclone of HepG2). Dimethyl sulfoxide (DMSO) is used as the negative control. The response is measured via reporter RNAs that are produced proportionately to the activity of corresponding transcription factor (here, PXR). This assay has been shown to be related to the body's response to toxic substances (Kliwer, Goodwin and Willson (2002)), so a higher response value for this assay endpoint can be interpreted as higher level of toxicity.

This work makes no monotonicity assumptions on the shape of the response. There is a nonzero baseline activity level of the nuclear pregnane X receptor in unstimulated hepatic cells, so both induction and suppression of the activity of this transcription factor are possible outcomes in response to stimulation, rendering a positive monotonicity constraint unsuitable. Furthermore, the lack of any monotonicity constraint (not specifically a positive one) leaves open the possibility for a biphasic response such as hormesis to be fit.

The tool used to quantitatively summarize a chemical's molecular structure is Mold2 (Hong et al. (2008)), which generates a set of 777 numeric descriptors using the simplified molecular-input line-entry system (SMILES) specification (Weininger (1988)); see Figure 3 for select Mold2 output for an example chemical, and Hong et al. (2012) for a discussion of the use of Mold2 in quantitative structure–activity relationship (QSAR) models.

In order to coherently model both structural and toxicological response variation, we propose a Bayesian partially Supervised Sparse and Smooth Factor Analysis (BS<sup>3</sup>FA) model. The model assumes structured variation in the molecular features  $x_i$  is driven by two sets of latent factors: call these  $F_{x\text{-specific}}$  and  $F_{\text{shared}}$ .  $F_{x\text{-specific}}$  is unrelated to the toxicological response and is responsible for structured molecular variability that does not impact toxicity.  $F_{\text{shared}}$  is assumed to drive variation in the toxicological response  $y_i$  and thus is responsible for structured molecular variability that *does* impact toxicity. The directions spanned by these two sets of latent factors can be thought of as the “toxicity-irrelevant” and “toxicity-relevant” spaces, respectively.

Chemical similarity can be characterized by proximity in this latent toxicity-relevant space, enabling a measure of distance with uncertainty quantification that is adapted to the

particular response space of interest. Such a metric is powerful because: (1) it is based on a subspace driving variation in activity, whereas proximity with respect to the full set of molecular descriptors does not necessarily mean proximity with respect to activity (Martin, Kofron and Traphagen (2002), Nikolova and Jaworska (2003)) and (2) it is purely statistically derived, requiring no knowledge of the fundamental chemical and biological processes responsible for the activity, as such information is not always available. Such an activity-relevant distance metric could be used by toxicologists in the design of diverse chemical libraries or to select new compounds to augment a screening collection such as ToxCast.

As with function-on-scalars regression approaches, the BS<sup>3</sup>FA model allows for the prediction of activity profiles for chemicals that have not yet been screened in ToxCast. It does so by embedding the full set of molecular features for a new chemical into the latent toxicity-relevant feature space  $F_{\text{shared}}$  and then projecting this embedding out to the activity space. The predicted dose-response profiles can be used to generate point and interval estimates for common univariate toxicological outcomes of interest, such as 50% activity concentration (AC50), maximum activity or the area under curve (AUC), which can be used in place of the as of yet unobserved *in vitro* results for that chemical.

The rest of the paper is organized as follows. First, we describe existing and potential approaches to modeling chemical structure and activity. Then, the BS<sup>3</sup>FA model is described, and its performance is compared to that of existing algorithms on simulated data sets. Next, a detailed analysis of the motivating application data set is considered, where the BS<sup>3</sup>FA model is run with Mold2 chemical features and the Attagene PXR assay from the ToxCast data set as input data. Finally, the results are discussed and future areas of research are highlighted.

## 2. Background.

QSAR models (see Figure 4) are based on the assumption that chemicals with similar features are likely to have similar effects. The ToxCast data poses two main challenges for QSAR modeling. First, it is often not trivial to characterize similarity in activity-relevant chemical feature space well (Martin, Kofron and Traphagen (2002), Nikolova and Jaworska (2003)). Second, the majority of QSAR models aim to relate structure to a *summary* of the data across times/doses (e.g., Liu et al. (2011), Patel et al. (2014), O'Connell and Lock (2019)) rather than to the *full* dose response curves.

We discuss these considerations in the context of existing QSAR models and other related approaches not yet applied to QSAR models. To the authors' knowledge, no existing approaches are able to address the challenge of learning a low-dimensional representation for multivariate feature data  $x_j$  partially supervised by sparse functional data  $y_j$ .

### 2.1. QSAR approaches for dose response profiles.

Two existing QSAR approaches have attempted to relate molecular descriptors to full dose response curves (Low-Kam et al. (2015), Wheeler (2019)). In Low-Kam et al. (2015), a Bayesian regression tree is defined over functions where each leaf represents a different

dose-response surface. This method was used to learn about the relationship between chemical properties and observed dose-response. However, the model lacks the ability to scale to the numbers of chemicals and molecular descriptors considered here. Furthermore, predictive performance was found to be lacking in leave-one-out analysis. Finally, the code was designed under the assumption that each chemical would be tested at the same doses with the same number of replicates at each dose.

The Bayesian additive adaptive basis tensor product (BAABTP) model (Wheeler (2019)) is designed purely for prediction. It learns basis functions via independent Gaussian process (GP) priors over the molecular structure space and the dose space. In the model, step one is to perform PCA on the set of Mold2 chemical descriptors. Step two is to use the principal feature space explaining 95% of the variation in this Mold2 descriptor set as the input to the distance kernel for the molecular structure GPs.

The BAABTP model has two major problems, both stemming from the Gaussian process prior over chemical structure. First, the model becomes computationally intractable when the number of chemicals increases past a few thousand—the number of chemicals tested in the ATG PXR assay has increased from under 1000 up to nearly 4000 in the time since the data were analyzed in Wheeler (2019), and this number will only continue to grow. Second, the GP priors over chemical structure rely on a concept of molecular distance based on total, rather than toxicity-relevant, variability. Mold2 descriptors are a numeric representation of the 2D structure of a chemical. Thus, while the leading principal components (PCs) account for the majority of structural variability across chemicals, these leading PCs may not be those most relevant to toxicity. Figure 5 illustrates this phenomenon, showing that chemicals may be close in PCA-based structure space but distant in response space. In other words, similarity in directions of highest structural variability does not necessarily correspond to similarity in directions of highest activity variability.

## 2.2. Toward the proposed approach.

Supervised and sparse functional PCA (supSFPCA) (Li, Shen and Huang (2016)) defines a hierarchical model to provide supervision for dimension reduction of functional data by another multi-output data source. A small modification to the penalty term used in this algorithm would allow for the opposite relation (i.e., to supervise the dimension reduction of the feature data  $x_j$  by functional dose-response data  $y_j$ ), but the larger issue is that the supSFPCA algorithm is designed such that the algorithm finds directions that maximize *unexplained* variability in  $x_j$  (i.e., what isn't accounted for by  $y_j$ ) rather than finding directions that maximize variability *explained* by  $y_j$ . As such, this algorithm is of little use for prediction or for learning about directions of variability of most relevance for toxicity.

The Bayesian latent factor regression model (LFRM) of Montagna et al. (2012) characterizes  $y_j$  as a linear combination of a high-dimensional set of basis functions. The basis functions themselves are fixed rather than learned which means that a large number of basis functions must be included in order to flexibly model many possible functional shapes. This high number of bases drives the choice to use a latent factor regression model on the basis coefficients rather than defining a regression model on the basis coefficients themselves. Thus, the latent factor scores enter the model in a “lower” layer in Montagna

et al. (2012) than they do in BS<sup>3</sup>FA and are a linear *function of* covariates  $x_i$  rather than being assumed to *drive variability in*  $x_i$ . A shrinkage prior is placed on the regression terms associated with the latent factor scores in the model, allowing covariates to impact certain facets of the functional response while others are shrunk to zero.

Our approach was inspired by the conceptual goal of separating variability shared by  $x_i$  and  $y_i$  into toxicity-relevant, toxicity-irrelevant and noise components, similar to the idea behind the joint and individual variation explained (JIVE) method (Lock et al. (2013)), and an analogous joint Bayesian factor model (JBFM) (Ray et al. (2014)). We utilize similar shrinkage tools as LFRM and the Bayesian function-on-scalars regression (B-FOSR) of (Kowal and Bourgeois (2020)), implementing sparsity-inducing coefficient-level priors to account for the high-dimensional predictors and imposing ordered sparsity so as to reduce the impact of the choice of latent subspace dimension. Unlike LFRM and like B-FOSR, the BS<sup>3</sup>FA model learns a flexible basis for the functional response rather than prescribing a set of basis functions. Like LFRM, JBFM, and B-FOSR, modeling takes place within a Bayesian framework for unified parameter estimation, prediction and uncertainty quantification about posterior summaries of interest, and model components are identifiable.

Like many of the approaches described in the previous sections, our model is able to predict the dose-response profiles for new chemicals. Unlike previous models, our model is able to learn a toxicity-relevant subspace underlying variation in both the chemical structure and toxicological response. This subspace can be used to describe a statistically-driven activity-relevant distance between chemicals. It can also provide insight into how toxicity-relevant variability manifests across both molecular structure and dose-response profiles.

### 3. BS<sup>3</sup>FA model.

Figure 6 gives a visual representation of the BS<sup>3</sup>FA model. BS<sup>3</sup>FA is able to: (1) learn a linear low-dimensional latent space underlying both molecular structure and activity, (2) develop a distance metric for chemicals relying only on molecular structures relevant to toxicity, (3) handle responses observed at a sparse, irregularly spaced set of doses and (4) enable activity predictions for chemicals having no observed dose-response information.

#### 3.1. Model specification.

**The characteristics of chemical structure and toxicological response can likely be summarized using fewer descriptors.**—As shown in Figure 1, many molecular descriptors tend to exhibit high correlation. Furthermore, their realized value can often be attributed to some underlying trait of the molecule (e.g., many descriptors are largely driven by molecule size; the descriptors number of carbon, number of oxygen, molecular weight and number of atoms in the molecule all have pairwise correlation above 0.9). Similarly, dose-response profiles exhibit a somewhat limited range of shapes, suggesting possibly a few underlying functional “building blocks” comprising variation in activity.

Factor modeling is a means by which to model variability in high-dimensional data via an underlying lower dimensional subspace. For some set of observations  $\{z_i\}_{i=1}^N$ , where  $z_i$  is

the  $P$ -dimensional vector of measurements for observation  $i$ , the traditional (nonjoint) factor model is

$$\begin{aligned} \mathbf{z}_i &= \Lambda \boldsymbol{\eta}_i + \boldsymbol{\epsilon}_i, \quad \boldsymbol{\eta}_i \sim \mathcal{N}_K(0, I), \quad \boldsymbol{\epsilon}_i \sim \mathcal{N}_P(0, \Sigma_0), \\ \Sigma_0 &= \text{diag}(\sigma_1^2, \dots, \sigma_P^2), \quad i = 1, \dots, N. \end{aligned} \quad (1)$$

The prior induced on the latent  $\mathbf{z}_i$  by integrating out the unknown  $\boldsymbol{\eta}_i$  is then

$$\mathbf{z}_i \sim \mathcal{N}(0, \Lambda \Lambda' + \Sigma_0), \quad (2)$$

yielding a lower dimensional representation of the covariance between measurements.

**Chemical features are often nonnormal (e.g., count, skewed continuous or binary).**—Many chemical descriptors from Mold2 are counts of particular elements (number of carbon, number of oxygen, etc.).

In order to allow this framework to encompass data of mixed type, define

$$x_{is} = f_s(z_{is}), \quad i = 1, \dots, N, \quad s = 1, \dots, S. \quad (3)$$

The particular link function  $f_s$  depends on the feature specification, allowing for mixed scale data via selection of an appropriate link by scale and type. Let  $f_s(z_{is}) = z_{is}$  or  $f_s(z_{is}) = \log(z_{is})$  for continuous  $x_{is}$ , with the latter chosen for strictly positive and positively skewed cases. Let  $f_s(z_{is}) = 1(z_{is} > 0)$  for binary  $x_{is}$ , where  $1(\cdot)$  is an indicator function taking the value of 1 when the argument is true and 0 when the argument is false. Categorical variables may be incorporated under this framework by transforming the  $C$  categories into  $C - 1$  binary variables indicating whether or not the categorical value for that individual took a given nonbaseline category value; the result is that either none or one of the  $C - 1$  variables will take on a value of 1. Finally, for count  $x_{is}$ , which may or may not be zero-inflated, let  $f_s(z_{is})$  be a rounding operator such that  $f_s(z_{is}) = 0$  if  $z_{is} < 0$  and  $f_s(z_{is}) = t$  if  $t - 1 < z_{is} < t$ , as specified in Canale and Dunson (2013).

**There is likely a shared low-dimensional space underlying chemical features and activity.**—BS<sup>3</sup>FA assumes that some underlying factors explain all of the variation in the dose-response curves and, jointly, part of the variation in the associated chemical features. Recall  $\mathbf{z}_i$  and  $\mathbf{y}_i$  denote  $S$ - and  $D$ -dimensional (latent) continuous features and observed dose-response curves, respectively, for observation  $i$ . Assume that the indexing of  $\mathbf{y}_i$  is such that the function is “in order” (for the ToxCast data, in order means that the  $D$  unique doses are sorted such that the doses increase with index). Also, for notational convenience assume that functional data  $\mathbf{y}_i$  are only observed once per index (the case of notationally awkward multiply observed doses is explicitly addressed in the Gibbs sampler described in the Appendix). BS<sup>3</sup>FA models

$$\begin{aligned}
 \mathbf{z}_i &= \boldsymbol{\mu}^z + \Theta \boldsymbol{\eta}_i + \Xi \mathbf{v}_i + \mathbf{e}_i, \\
 \mathbf{y}_i &= \boldsymbol{\mu}^y + \Lambda \boldsymbol{\eta}_i + \boldsymbol{\epsilon}_i, \\
 i &= 1, \dots, N.
 \end{aligned} \tag{4}$$

The form of the above model is that of a set of linked factor models. Note that, although there are three non-error component pieces aside from the global mean terms (namely,  $\Theta \boldsymbol{\eta}_i$ ,  $\Xi \mathbf{v}_i$ , and  $\Lambda \boldsymbol{\eta}_i$ ), there are only two unique factor vectors:  $\boldsymbol{\eta}_i$  and  $\mathbf{v}_i$ . These factors are highly interpretable. The term  $\boldsymbol{\eta}_i$  represents the shared latent space underlying structured variability in both  $\mathbf{z}_i$  and  $\mathbf{y}_i$  (note that it appears in both factor models). In the expression for  $\mathbf{y}_i$ , it is the sole factor vector and in that for  $\mathbf{z}_i$  it is one of two factor vectors. Thus, it is responsible for all structured variation in  $\mathbf{y}_i$  but only part of the structured variation in  $\mathbf{z}_i$ . The term  $\mathbf{v}_i$  represents structured variation in  $\mathbf{z}_i$  that is *unrelated* to  $\mathbf{y}_i$ .

By not including unique response-only factors (i.e., factors not assumed to underlay chemical structure) in the mean formulation for  $\mathbf{y}_i$ , the model assumes that the mean dose-response curves can be constructed from factors that are all also present in the chemical feature data. Making this assumption allows for easier identification and accurate estimation of parameters associated with chemical features predictive of activity, which is the main goal of this work. The trade-off of this assumption is that if the feature set is inappropriate (e.g., not containing enough information by which to estimate activity) and/or the dose-response profiles are truly driven by some individual factors not related to chemical features, then the estimates for new chemicals may be overconfident and the model fit may be poor. The shrinkage prior on  $\Theta$ , discussed in a later section, offers some mitigation of this risk.

If  $\mathbf{z}_i$  and  $\mathbf{y}_i$  are mean-centered prior to analysis, fix  $\boldsymbol{\mu}^z$  and  $\boldsymbol{\mu}^y$  to be zero vectors for computational savings. In practice, mean centering is not sensible to perform for count or binary features (i.e., for  $s$  s.t.  $x_{is}$  is binary or count). For noncentered  $x_{is}$ ,  $\mu_s^z$  is given a Cauchy prior expressed as  $\mu_s^z \sim \mathcal{N}(0, \zeta_s^{-1})$ ,  $\zeta_s \sim \text{Ga}(0.5, 0.5)$ . Since the Cauchy distribution has high density around zero and heavy tails, it is able to capture meaningful signals while still encouraging shrinkage. If the  $\mathbf{y}_i$  are not mean-centered,  $\boldsymbol{\mu}^y$  is given a GP prior analogous to the prior placed on the columns of  $\Lambda$ , discussed in the next section.

In the above model the mean of  $\mathbf{z}_i$  is  $\boldsymbol{\mu}^z + \Theta \boldsymbol{\eta}_i + \Xi \mathbf{v}_i$  and  $\mathbf{e}_i$  is a term for unstructured noise in  $\mathbf{z}_i$ . Similarly, the mean of  $\mathbf{y}_i$  is  $\boldsymbol{\mu}^y + \Lambda \boldsymbol{\eta}_i$  and  $\boldsymbol{\epsilon}_i$  is the unstructured noise term for  $\mathbf{y}_i$ . The priors on the shared factors  $\{\boldsymbol{\eta}_i\}$ , the  $X$ -specific factors  $\{\mathbf{v}_i\}$  and error terms are set to be those typically used in factor analysis,

$$\begin{aligned}
 \boldsymbol{\eta}_i &\sim \mathcal{N}_K(0, I), & \mathbf{v}_i &\sim \mathcal{N}_J(0, I), \\
 \mathbf{e}_i &\sim \mathcal{N}_S(0, \Sigma_X), & \Sigma_X &= \text{diag}(\sigma_{X,1}^2, \dots, \sigma_{X,S}^2), \\
 \boldsymbol{\epsilon}_i &\sim \mathcal{N}_D(0, \Sigma_Y), & \Sigma_Y &= \text{diag}(\sigma_Y^2, \dots, \sigma_Y^2).
 \end{aligned} \tag{5}$$



Homoscedastic variance is assumed for dose-response curves  $Y$ . Fix  $\sigma_{X,s}^2$  to 1, if feature  $s$  is binary, for reasons of identifiability.

Note that, as of yet, we have not fully addressed how to structure the factor model in light of the functional nature of the  $y_j$ , nor the issue of many entries in  $z_j$  likely being unrelated to  $y_j$ . The following sections will describe how structure can be imposed on  $\Lambda$  and  $\Theta$ , respectively, in light of these considerations.

**The dose-response curve data are functional in nature.**—Figure 2 shows the noisy observations from a set of example chemicals, but the underlying signal represents a smooth curve relating chemical dose to response.

For functional data it is preferable to have each loading vector (i.e., each column of  $\Lambda$ ) itself be functional. The desired smoothness of the mean curves underlying noisy observations  $y_j$  can thus be imposed via the choice of smooth priors on the loading matrix  $\Lambda$ . Let  $\lambda_k$  denote the  $k$ th column of  $\Lambda$ , so

$$\Lambda = \begin{bmatrix} | & | & \dots & | \\ \lambda_1 & \lambda_2 & \dots & \lambda_K \\ | & | & \dots & | \end{bmatrix}.$$

In order to learn rather than prescribe smooth bases, columns of  $\Lambda$  are modeled as  $D$ -dimensional Gaussian processes,

$$\lambda_k \sim \mathcal{GP}(0, c_k(\cdot)), \quad c_k(d, d') = \alpha_k^2 e^{-\frac{(d-d')^2}{2\ell^2}}, \quad k = 1, \dots, K. \quad (6)$$

The Gaussian process function variance  $\alpha_k^2$  is comprised of two components: a global inverse variance term  $\phi$  and a column-specific inverse variance term  $\tau_k$ . The column-specific inverse variance term  $\tau_k$  utilizes the multiplicative gamma process prior of Bhattacharya and Dunson (2011), leading to stochastic shrinkage of columns of  $\Lambda$  toward 0 by index:

$$\alpha_k^2 = (\phi \tau_k)^{-1}, \quad \tau_k = \prod_{h=1}^k \delta_h, \quad k = 1, \dots, K, \quad (7)$$

$$\phi \sim \text{Ga}(g_\phi/2, g_\phi/2), \quad \delta_1 \sim \text{Ga}(a_1, 1), \quad \delta_h \sim \text{Ga}(a_2, 1), \quad h \geq 2.$$

Following the note by Durante (2017) on hyperparameter selection, set  $a_1 = 2.1$  and  $a_2 = 3.1$  in equation (7). The value of  $g_\phi$ , the hyperparameter for the global function precision of the GP, should be chosen to reflect the scale of the data.

This stochastic shrinkage leads to an effective truncation of the factors and an automatically learned dimension of the latent space, so long as  $K$  is chosen large enough (whether  $K$  is adequately large can be assessed by monitoring the convergence of  $\alpha_k^2$  to 0 as  $k$  approaches  $K$ ). To see why, consider the model for  $y_j$  written in expanded form:  $y_j = \lambda_1 \eta_{j,1} + \lambda_2 \eta_{j,2} + \dots + \lambda_K \eta_{j,K} + \epsilon_j$ . Assuming  $K$  is large enough, as  $k$  approaches  $K$ , the vector  $\lambda_k$  should be

approximately the 0-vector, meaning that many of the later terms will contribute negligibly to the mean of  $y_j$ ; see Figure 7 for a visualization of  $\Lambda$  and column shrinkage.

The linked nature of the factor models leads to an induced prior on the covariance between elements of  $z_j$  and elements of  $y_j$ ; specifically,  $\text{Cov}(z_{is}, y_j(d)) = \sum_k \theta_{sk} \lambda_k(d)$ . The GP prior on columns of  $\Lambda$  implies that the covariance between a given feature and a given point on the dose-response profile is a smooth function of dose, a desirable model characteristic.

**Many features are likely unimportant for certain aspects of chemical activity.**

—That is, if a set of chemical descriptors has not been carefully selected to be toxicity-relevant, it is unlikely that all are related to the shape of the dose-response curves. Even if all features are toxicity-relevant, it is plausible that features will impact different pieces of the toxicity profile (e.g., some features may impact the steepness of the dose-response, while others may impact the height of the final plateau). The way to encourage such a relationship is via elementwise shrinkage (i.e., zeros in entries) of the factor loadings matrix  $\Theta$ .

Shrinkage on elements  $\theta_{sk}$  of  $\Theta$  is desirable because it is likely that for a given factor many features have negligible impacts on the associated component of the functional  $y_j$ . Explicitly,

$$\Theta = \begin{bmatrix} \theta_{11} & \theta_{12} & \theta_{13} & \dots & \theta_{1K} \\ \theta_{21} & \theta_{22} & \theta_{23} & \dots & \theta_{2K} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \theta_{S1} & \theta_{S2} & \theta_{S3} & \dots & \theta_{SK} \end{bmatrix}.$$

There is a very rich literature proposing elaborate shrinkage and sparsity priors for factor loadings (e.g., Knowles and Ghahramani (2011), Meng et al. (2010), Pati et al. (2014), Yoshida and West (2010)). We opt for a horseshoe prior (Carvalho, Polson and Scott (2010)) modified for simple sampling (Makalic and Schmidt (2016)) on entries  $\theta_{sk}$  of  $\Theta$ ,

$$\begin{aligned} \theta_{sk} &\sim \mathcal{N}(0, \beta^2 \gamma_{sk}^2 \tau_k^{-1}), \\ \beta^2 &| t \sim \text{IG}(1/2, 1/t) \\ \gamma_{sk}^2 &| b_{sk} \sim \text{IG}(1/2, 1/b_{sk}) \\ \{b_{sk}\}, t &\sim \text{IG}(1/2, 1), \\ s &= 1, \dots, S, \quad k = 1, \dots, K. \end{aligned} \tag{8}$$

The horseshoe component of this prior is in the hierarchical hyperprior on global variance term  $\beta^2$  and local variance term  $\gamma_{sk}^2$ . The column-specific variance  $\tau_k^{-1}$  applies stochastically increasing shrinkage as column index increases; see Figure 8 for a visualization of  $\Theta$  and column shrinkage.

Recall  $\text{Cov}(z_{is}, y_j(d)) = \sum_k \theta_{sk} \lambda_k(d)$ . The shrinkage prior on entries of  $\Theta$  allows for the possibility that a given chemical feature may not have a contribution from some or any of the latent factors (i.e., that  $\theta_{sk}$  for such feature/factor combinations is near 0). In this case the covariance between the feature and points along the dose-response profile will not depend

on those factor(s). In the extreme case that the  $s$ th feature is unrelated to toxicity,  $\theta_{sk}$  should be shrunk toward zero for all  $k$ , leading  $z_{is}$  and  $y_i(d)$  to be uncorrelated.

**The dimension of the latent space underlying the dose-response curves and of that underlying toxicity-relevant features should be the same.**—The same column-specific precision  $\tau_k^{-1}$  is used for both  $\Theta$  and  $\Lambda$ , so relative column shrinkage is applied consistently across the two matrices. The result is the same effective truncation on the number of latent factors in the joint space. This specification, along with the common  $\eta_i$  is what allows the shared directions of variability between  $y_j$  and  $x_j$  to be learned. Distance can be defined over the  $\eta$ -vector with elements  $\eta_k$  weighted by precision  $\tau_k^{-1}$  to give a sense of closeness in “ $\eta$  space” that reflects the true amount of information contained in each latent direction.

**There is additional variability in chemical structure beyond that which impacts chemical activity.**—Unless all chemical features were carefully hand-selected to be toxicity relevant for the specific assay considered in the model (a tall task and unlikely to be completely true no matter how careful the selection), accounting for variability in  $z_j$  shared with  $y_j$  will not capture all structured variability in chemical features.

After accounting for the variability in  $z_j$  shared with  $y_j$  (via latent factor  $\eta_j$ ),  $z_j$  may still have structured variation due to individual latent factor  $\nu_j$ . Let  $z_i^* = z_i - \theta\eta_i$ . Then,  $z_i^* = \Xi v_i + e_i$  which once again looks like a traditional factor model. The direct application of priors in Bhattacharya and Dunson (2011) is used for elements of  $\Xi$ . Specifically, elements  $\xi_{s,j}$  of  $\Xi$  are given prior

$$\begin{aligned} \xi_{s,j} \mid \kappa_{sj}, \omega_j &\sim \mathcal{N}(0, \kappa_{sj}^{-1} \omega_j^{-1}), \quad \omega_j = \prod_{h=1}^j \zeta_h, \quad s = 1, \dots, S, j = 1, \dots, J, \\ \kappa_{sj} &\sim \text{Ga}(g_k/2, g_k/2), \quad \zeta_1 \sim \text{Ga}(m_1, 1), \quad \zeta_h \sim \text{Ga}(m_2, 1), \quad h \geq 2. \end{aligned} \quad (9)$$

Stochastic column-specific shrinkage via the  $\omega_j^{-1}$  term removes the need to select an ideal number of factors  $J$  and allows for simply selecting  $J$  “large enough.” This formulation also allows for efficient Gibbs sampling of the posterior. Following the note by Durante (2017) on hyperparameter selection, set  $m_1 = 2.1$  and  $m_2 = 3.1$  in equation (9). The value of  $g_k$ , the hyperparameter for the entry-level precision terms of  $\Xi$ , should be chosen to reflect the scale of the data.

If there is in fact, no additional variability in  $z_j$  beyond that shared with  $y_j$ , the following specification allows for all columns of  $\Xi$  to be shrunk to 0-vectors. This case reduces to a fully joint factor model in which all variability in  $z_j$  is shared with  $y_j$ .

**Chemical activity is not necessarily measured on a fully observed, regularly spaced grid.**—There are a handful of common dose measurements at which the majority of chemicals are measured (see Figure 11), but there are many chemicals whose observations are less regular. Furthermore, some chemicals have multiple observed dose-response curves.

The issue of irregular spacing between the unique values associated with the indices is handled automatically via the use of GPs for modeling columns of  $\Lambda$ . The covariance between points is defined by the kernel for any pair of input values (see equation (6)) and not dependent on a regular measurement grid.

**Toxicologists may wish to report different components and/or summaries of predicted dose-response curves.**—For example, they may be interested in the dose value at which the response first exceeds some threshold, the maximum response value reached, the area under the curve, etc. Each of these summaries provides different information about the dose-response relationship. An advantage of a Bayesian formulation is that we can obtain posterior samples for any functional of the dose-response curve trivially, with these samples then used to obtain point and interval estimates.

### 3.2. Posterior computation.

The posterior for the BS<sup>3</sup>FA model is not available in closed form. However, closed form full conditional distributions of the parameters associated with the model allow the use of a straightforward Gibbs sampler for these draws. Samples obtained directly from this Markov chain Monte Carlo (MCMC) algorithm allow for the calculation of posterior means, simultaneous bands (Meyer et al. (2015)), and credible intervals for identifiable model components, including the predicted mean, covariance, and noise variance of  $X$  and  $Y$ . A post-processing step to resolve rotational ambiguity and account for label/sign switching allows for identifiability of the individual model components, including the factor scores  $\eta$  and loadings  $\Lambda$  and  $\Theta$  (code modified from [https://github.com/poworoznek/sparse\\_bayesian\\_infinite\\_factor\\_models](https://github.com/poworoznek/sparse_bayesian_infinite_factor_models)). Full details on the Gibbs sampler steps and initialization are included in the Supplementary Material (Moran et al. (2021)).

### 3.3. Code base and reproducibility.

Code for simulating data and sampling from the BS<sup>3</sup>FA model, along with a user manual, are made available at <https://github.com/kelrenmor/bs3fa>. A hands-on demonstration of the package is available at <https://www.youtube.com/watch?v=qLyxBQ-sVcY>. Code specific to this paper (i.e., to reproduce the simulations, figures, and results) is provided in the Supplementary Material (Moran et al. (2021)).

### 3.4. Simulation study.

Simulation studies were performed in order to assess the ability of BS<sup>3</sup>FA to learn the true toxicity-relevant distance between chemicals, its predictive performance and the model fit. Two broad categories of simulations were performed: first, those in which the true data generating process aligns with model assumptions (i.e., when data are simulated from a partially shared latent factor model) and when it does not (i.e., when data are simulated from something other than a factor model). In the former category we also assess how well model subcomponents can be learned.

For all simulations, 25% of simulated “chemicals” are held out. That is, rather than withholding 25% of dose-response observations across chemicals, we withhold all dose-response data for each hold-out chemical. For distance performance, BS<sup>3</sup>FA is compared to Euclidean

distance using all features, using the features selected in the B-FOSR model from the *foser* R package via an extension of the decision analytic approach of Hahn and Carvalho (2015) to the functional data setting (Kowal and Bourgeois (2020)), using the principal component scores explaining 95% of the variability in the data found in PCA, and using the selected variables from a frequentist FOSR using the *foser.vs()* function (FOSR-VS) from the *refund* R package with the PC scores used as inputs. PC scores were used instead of the full set of features for FOSR-VS because the *foser.vs()* function returned an error saying the dimension of  $y_i$  was not high enough relative to  $x_i$  (i.e., that  $D$  was too small relative to  $S$ ) when the full feature set was used. The performance of each method is compared by computing the correlation between predicted and true distance, a choice made because the scale of distance varies across methods and, by virtue of the structure of the simulated data, we don't expect any extreme outliers in this space. For predictive performance, BS<sup>3</sup>FA is compared to the BAABTP model (Wheeler (2019)), B-FOSR, FOSR-VS, and least absolute shrinkage and selection operator (LASSO) using each covariate, dose and all pairwise interactions.

For simulations in which the true data generating process aligns with model assumptions, the true dimension of the latent toxicity-relevant space  $K$  was varied, taking values 1, 3 and 5. For each  $K$ , the true dimension of the latent toxicity-irrelevant space  $J$  was varied from 0 to 20 in intervals of five. At each combination of  $K$  and  $J$ , 100 data sets were simulated with  $N=300$ ,  $D=10$ , and  $S=40$ . For roughly half of the chemicals in each data set,  $\eta_i$  was set to a zero vector for that chemical (i.e., for each simulated data set there was a 50% chance that any given chemical was inactive). Further simulation details are provided in the Supplementary Material (Moran et al. (2021)). Overall, the model does quite well at capturing the structure of the noise variance and the true components  $\Lambda$  and  $\Theta$ .

Figure 9 shows the correlation between entries in the true pairwise distance matrix (i.e., the Euclidean distance between true latent factors  $\eta$ ) and the predicted pairwise distance matrix for holdout chemicals. We see that even in the case of small  $J$ , performing PCA on the  $S$ -dimensional  $X$  matrix obscures the true distance in the latent space. As  $J$  increases, the correlation between the chemical distance in the true  $\eta$  and that in either PCA space or Euclidean space drops quickly. This phenomenon also occurs for distances computed via variable selection using B-FOSR or FOSR-VS. The BS<sup>3</sup>FA model has stable high correlation across all values of  $J$  and  $K$ .

Figure 10 shows the mean squared predictive error (MSPE) for the simulated hold-out chemicals' dose-response mean functions. Although the performance of all models deteriorates as the amount of "superfluous" information in  $X$  increases (i.e., as  $J$  increases), the BS<sup>3</sup>FA and B-FOSR models are the most robust, showing superior performance across all values of  $K$  and  $J$ . The BS<sup>3</sup>FA consistently outperforms the B-FOSR model, with MSPE of the latter on average 1.5 times higher than that of the former (the 95% quantile of this multiplicative factor ranges from 1.1 to 2.4 across all simulations and values of  $K$  and  $J$ ). The BAABTP model appears most sensitive to the value of  $J$ , with MSPE near that of BS<sup>3</sup>FA model when  $J$  is small, but among the worst MSPE when  $J$  is high. The LASSO model is able to perform fairly well when  $K$  is small, but, as  $K$  increases, it is unable to

learn the more complicated relationship between  $x_j$  and  $y_j$ . The FOSR-VS model becomes increasingly sensitive to  $J$  as  $K$  increases.

Coverage of the simultaneous bands for the BS<sup>3</sup>FA model remains close to nominal across all values of  $K$  and  $J$ . Additionally, these band widths are narrowest subject to (at least) nominal coverage across competitors and for all values of  $K$  and  $J$ . The coverage of the BAABTP model is consistently below nominal and decreases as  $J$  increases, another reflection of its overall poor predictive performance. The coverage of the B-FOSR model is much higher than nominal, nearing 1 across all values of  $K$  and  $J$ , and the simultaneous bands are wider on average than those of BS<sup>3</sup>FA. Detailed results for both coverage and band width are available in the Supplementary Material (Moran et al. (2021)).

In the ToxCast analysis following this section, we discuss results from one possible method of deeming a chemical active. Namely, call a chemical active if its global Bayesian p-value (Meyer et al. (2015)) is less than 0.05. The null hypothesis in this case is that the dose-response mean function equals 0 everywhere, that is, that a chemical is inactive. If alternatively activation (suppression) is of specific interest, one could use the additional requirement that the direction of observed effect be positive (negative) in order to reject the null. The true positive rate (TPR), false positive rate (FPR) and false discovery rate (FDR) of this method on the simulated data are shown in Table 1. Across all values of  $K$  and  $J$ , the TPR is generally high, and the FPR and FDR are near zero. Note that the TPR increases with  $K$  and decreases with  $J$ , while FPR and FDR remain stable across values of  $K$  and  $J$ . That is, the model seems more likely to identify chemicals as active, as the dimension of the latent toxicity-relevant space increases, but loses sensitivity when there is more toxicity-irrelevant information. Results for the B-FOSR method, provided in the Supplementary Material (Moran et al. (2021)), show slightly improved TPRs but at the cost of much higher FPRs and FDRs (whereas for BS<sup>3</sup>FA the FPR and FDR are consistently near zero, these values range from 0.1 to 0.5 across  $K$  and  $J$  for B-FOSR).

When there is misalignment between the structure assumed by the BS<sup>3</sup>FA model and the true data generating process, BS<sup>3</sup>FA is still able to predict similarly to or better than the competitors. As with the well-aligned simulation, BS<sup>3</sup>FA is robust to increasing “superfluous” information in  $X$ . A similar story can be seen in the coverage and distance results for the misaligned simulation. Even when the assumed latent factor model is not the model from which data are simulated, the coverage of BS<sup>3</sup>FA is close to nominal. B-FOSR, on the other hand, suffers from much higher than nominal coverage across all simulations while BAABTP suffers from much lower than nominal coverage as “superfluous” information in  $X$  increases. BS<sup>3</sup>FA is still able to recover a distance metric that is highly correlated (mean correlation 0.97, range 0.94 to 0.99 across simulations) with the distance in the true relevant  $X$  dimensions, although B-FOSR more perfectly learns such a metric. Visual results are shown in the Supplementary Material (Moran et al. (2021)).

#### 4. Relating chemical structure to toxicological response.

Data preprocessing steps and results of the analysis of the ToxCast ATG PXR assay are discussed in the following subsections. The structure of BS<sup>3</sup>FA allows for learning about

structured variability in both the feature and response space, prioritizing chemicals for future evaluation, and predicting chemical activity for as yet unobserved chemicals.

#### 4.1. ToxCast setup.

Observations below the cytotoxicity limit for 3,540 Phase 1, Phase 2 and e1k chemicals tested in the AttaGene PXR assay are included in our data analysis. The structure information for each chemical is summarized by 777 Mold2 chemical features (Hong et al. (2008)). As discussed previously, BS<sup>3</sup>FA has the advantage of being able to effectively ignore toxicity-irrelevant features via shrinkage on elements of  $\Theta$ , making the careful curation of a feature set unnecessary.

Chemicals having no provided SMILES information ( $n = 405$ ) were omitted from further analysis. The result is 3,135 chemicals having Mold2 descriptions of their chemical structure. Note that some chemicals have identical Mold2 output. These sets of chemicals are, in effect, considered a single chemical (i.e., treated as multiply observed dose-response curves by the model). For analysis the number of “unique” chemical sets (i.e., the number of unique SMILES represented across the 3,135 chemicals) is  $N = 3,070$ . The eight most common dose concentrations are  $-1.05, -0.52, -0.1, 0.3, 0.85, 1.3, 1.85, 2.3$  log uM (see the vertical white bands in Figure 11), but BS<sup>3</sup>FA allows both common and unique doses. As the bulk of chemicals have no information about extremely low ( $< -2$  log uM) dose activity (see Figure 11), the data considered are the  $D = 38$  unique doses  $-2$  log uM, out of the 56 total unique doses. Approximately 4% of chemicals have multiple observed dose-response curves, for example, Allethrin and Clorophene from Figure 2. While such repeated measurements may come from different experimental runs, unfortunately the data set does not provide information on which point is associated with which run, so we assume independence in this respect rather than, for example, adding a run-specific random effect term.

Mold2 is used to generate a set of 777 numeric molecular descriptors associated with each chemical. As noted above, some chemicals exactly shared these descriptors due to Mold2’s inability to capture certain differentiating structures (see the Supplementary Material (Moran et al. (2021)) for an example); these “identical” chemicals were treated as multiple observed chemicals. After removing features having no variability (99 total, including, e.g., features equalling zero for all chemicals, such as number of 11-membered rings and number of Argon), features with duplicated entries (16) or features having  $> 99\%$  of chemicals sharing a feature value (99, e.g., only one chemical has any aromatic group urea derivatives),  $S = 563$  features remain. As a further preprocessing step, the continuous variables are scaled to have mean 0 and variance 1. Further information on creating and using Mold2 descriptors is included in the Supplementary Material (Moran et al. (2021)).

In order to mimic the scenario of using the BS<sup>3</sup>FA model to prioritize chemicals for further screening, we hold out all of the dose-response observations for 25% of chemicals in the data set. That is, we provide the model with these chemicals’ structure but not their dose-response curves. Of interest will be how similar these unobserved chemicals are to known active chemicals and the AC50 (the dose at 50% of maximum activity) for unobserved chemicals predicted by the model to be activity-increasing. Recall for the purpose of this

analysis, we consider a chemical activity-increasing if the global Bayesian p-value (Meyer et al. (2015)) of its dose-response profile is less than 0.05 and the direction of observed effect is positive.

A global mean term  $\mu^Y$  is included in the model (see Figure 12) to account for a nonzero average profile rather than mean centering, as the low number of observations at less common doses leads to a “noisy” center. Additionally, the dose-response matrix is scaled by a multiplicative factor so that the scaled  $X$  and  $Y$  matrices have the same Frobenius norm (a measure of total variation); specifically, set  $Y = \|X\|_F \frac{Y}{\|Y\|_F}$ . This rescaling keeps larger matrices from dominating when learning the shared column-specific shrinkage terms  $\{\tau_k\}$  or the shared score vectors  $\{\eta_j\}$  (e.g., when  $S$  is much larger than  $D$ , as in this setting).

We ran the sampler for 40,000 iterations. After an initial burn-in of 20,000 iterations, every 10th sample was saved. Computation time was approximately eight hours on a 2016 Mac-Book pro with a 2.9 GHz Intel Core i7 processor. Trace plots of model predictions show good mixing; these, along with those of model components and an assessment of the sufficiency of the chosen  $K$  and  $J$  values, are available in the Supplementary Material (Moran et al. (2021)).

#### 4.2. Model components.

The learned matrix  $\Lambda$ , shown on the left side of Figure 13, provides a snapshot of the directions of structured variation present in the dose-response data. The first column of  $\Lambda$ , shown in the top middle of Figure 13, is the dominant factor loading (i.e., the factor loading having the largest estimated norm). Unsurprisingly, this vector takes the shape of a prototypical dose-response curve. Later columns of  $\Lambda$  act to provide smooth deviations from this prototypical shape. For example, the third column characterizes a flat profile until around  $-1 \log \text{uM}$  followed by a near-linear increase until leveling off at a dose value just under  $2 \log \text{uM}$ , while the fourth column shows an initial start point above zero followed by a gentle U-shaped dip.

The learned matrix  $\Theta$ , the values of which are shown in Figure 14, provides a snapshot of the directions of structured toxicity-relevant variation present in the feature data. The bulk of the estimated entries are very close to 0 due to the shrinkage effect of the horseshoe prior. The BS<sup>3</sup>FA model structure allows us to interpret the nonzero entries of a given column of  $\Theta$  as being those related to the particular structure present in the corresponding column of  $\Lambda$ . For example, the significantly nonzero entries of the first column of  $\Theta$  are those associated with the prototypical activity profile seen in the top row of the middle column of Figure 13. By absolute magnitude the largest such features include the number of group esters, the number of group X-C on aromatic ring, the sum of the topological distance between the vertices O and Cl, the number of Chlorine and the sum eigenvalue weighted by polarizability distance matrix.

The chemicals having extreme positive values of  $\eta_1$  will be those for which the dose-response profile has a large component due to  $\lambda_1$  and a large chunk of toxicity-relevant molecular variability described by the first column of  $\Theta$ . In the training set the chemicals



having the largest expected value for  $\eta_1$  are Calcium bromide, Nickel(II) chloride, Dibromomethane, Zinc chloride, and Dimethylamine. All are known toxins. Several of these chemicals have in common the presence of Cl, so it is unsurprising that features involving Cl appeared amongst the high-value loadings components for the first column of  $\Theta$ .

### 4.3. Distance learning.

Figure 15 shows the predicted pairwise distance matrix between a set of example chemicals chosen as clusters of chemicals in the training set closest to specific recognizable hold-out chemicals. Included are a cluster of similar low-activity chemicals (the training chemicals nearest to hold-out chemical Saccharin) in the bottom left block and a cluster of similar high-activity chemicals (the training chemicals nearest to hold-out chemical Clomiphene citrate (1:1)) in the central block. A handful of miscellaneous chemicals that are fairly isolated in  $\eta$  space relative to the other included chemicals are shown in the top-right block.

When selecting future chemicals for prioritization, one can either seek to “fill in” the space around chemicals of known toxicity relevance or to “venture out” into spaces not near any currently tested chemicals. While addressing this experimental design problem is outside the scope of this work, we assume for the sake of exposition that both possible avenues are of interest. We further assume that the set of hold-out chemicals represents the space of options for further *in vitro* testing.

Assuming the central cluster of chemicals in Figure 15 is of interest for more targeted exploration, we could select the hold-out chemicals closest to that set to test further. In terms of average distance between each cluster member, the four closest chemicals in the hold-out set are Clomiphene citrate (1:1) (i.e., our “seed” chemical for this training group for which the model predictive performance can be seen in Figure 17), 4-Hydroxytamoxifen, Tamoxifen, and Tamoxifen citrate. Research has suggested similarity in action between Clomiphene citrate and Tamoxifen (Dhaliwal et al. (2011), Seyedoshohadaei, Zandvakily and Shahgeibi (2012)), lending credence to this distance measure’s accounting of activity-relevant similarity.

If, on the other hand, our goal was to test new chemicals that are least similar to observed chemicals, then we could choose the hold-out chemicals having the largest minimum distance to a training chemical. Assuming we chose such chemicals iteratively, we would select Tetracosafuorotetradecahydrophenanthrene (a solvent used in the preparation of certain polymers), Strychnine hemisulphate salt (a pesticide used for rodent and bird control), and Cadmium dinitrate (a colorant and photographic flash powder component). That these chemicals are the most distant in  $\eta$  space from both the training set and each other suggests that the model considers them to have distinctive activity relevant variability. See Figure 16 for their structure diagrams; the predicted activity for Cadmium dinitrate is also shown in Figure 18.

### 4.4. Prediction.

Overall, the MSE between the data and the predicted dose-response profiles is 0.24 for the training chemicals and 0.30 for the hold-out chemicals. For comparison, the straw man

model using the mean of the training data at a given dose to predict the hold-out data at that dose leads to an MSE of 0.37 and predicting all hold-out dose-response profiles to be inactive (i.e., 0 everywhere) yields an MSE of 0.47. LASSO using dose, feature and the interaction between dose and feature gives an MSE of 0.31 for the training chemicals and 0.40 for the hold-out chemicals. The B-FOSR model, as run using the `fOSR( )` function in the `fOSR` package, had unstable predictions, as the amount of missingness in the training data was too high (recall Figure 11) for the imputation scheme used in the code—the MSE was in the 100s. Limiting the training data to only include the most commonly observed dose levels did not resolve this issue. The `fOSR.vs( )` function simply returned an error due to the amount of missingness in the training data. The BAABTP model was not run for the toxicity data for computation time reasons.

The 95% prediction intervals for the BS<sup>3</sup>FA model run cover 94.2% of the training data and 92.6% of the hold-out data, respectively. Unsurprisingly, hold-out chemicals having lower coverage also tend to have higher MSE. The coverage for specific hold-out chemicals is inversely related to the minimum distance between that chemical and its closest neighboring training chemical, while the MSE is directly related to the minimum distance between that chemical and its closest neighboring training chemical. That is, as a hold-out chemical moves farther away from other training chemicals, on average its coverage and MSE become worse. This finding suggests that increasing the amount of training data with specific care to the lesser-known regions of the chemical feature space would likely improve the model's MSE and coverage.

We say a chemical is predicted to be activating, that is, to increase activity, if the global Bayesian p-value (Meyer et al. (2015)) for the predicted dose-response profile is less than 0.05 and the 95% simultaneous bands exceed 0 for at least one dose value. Figures 17 and 18 show model-predicted mean dose-response (MDR) curves along with samples of the AC50 value, that is, the dose at which the dose-response curve is at half of its maximal value, for activating hold-out chemicals. The predicted mean and 95% posterior bands for the MDR curves are smooth due to the underlying structure of  $\Lambda$ . The proportion of hold-out chemicals deemed activating by our model is nearly twice as high among the population of chemicals that are heavily tested (i.e., that have 10 or more observations). Since chemicals that are known to have toxic effects tend to be more heavily tested, this finding is suggestive of the model's capability to detect activity.

Although the predictive ability of the model is imperfect (e.g., 1, 2, 3, 4, 5, 6-Hexachloro-cyclohexane in Figure 17 is under-predicted), overall performance appears reasonable. Particularly poorly predicted chemicals, examples of which are shown in the Supplementary Material (Moran et al. (2021)), tend to have shapes that differ from the common profiles and/or be farther in  $\eta$  space from training data than well-predicted curves. Also note that, although the observations above the reported cytotoxicity limit were removed prior to running the model, the hormesis shape (i.e., the downturn at the end of the predicted dose-response profile) remains prominent in highly active chemicals, because it is a feature of the first column of  $\Lambda$ , which has the highest column norm and drives large scale variation across profiles.

In the context of this experiment, we fail to reject the null hypothesis that a chemical is inactive if its global Bayesian p-value (Meyer et al. (2015)) is greater than 0.05 (i.e., if the bounds of the 95% posterior simultaneous bands for the predicted MDR curve include zero at all points). For convenience, we refer to these chemicals as being predicted to be inactive by the model. Figure 19 shows model-predicted inactive dose-response curves for hold-out chemicals. As before, the predictions are smooth and appear reasonable relative to the true data. On average, hold-out chemicals deemed inactive under the criteria outlined above have a lower true maximum observed response than those deemed activity-increasing.

The choice of how to prioritize chemicals for future evaluation is flexible. Assuming there are in fact no dose-response data for the hold-out chemicals, a simple scheme by which chemicals could be selected for *in vitro* study, based on their BS<sup>3</sup>FA predictions, would be to screen all chemicals for which the lower limit of the  $1 - \alpha$  posterior simultaneous bands for the MDR curve exceeds some threshold (e.g., 0). The value of  $\alpha$  could be selected with attention to the resources available, a larger  $\alpha$  would lead to more chemicals being screened, whereas a smaller  $\alpha$  would mean only those chemicals the model is most confident about would be screened. Once the set of chemicals are selected for further testing, the order of screening could be determined by chemicals' expected AC50 value, by the maximum value of their predicted MDR curves, by the highest value taken by the lower  $\alpha/2$  simultaneous band for expected response, by their proximity in latent space to known toxic chemicals, or by another metric of interest.

For example, if one chooses  $\alpha = 0.05$ , then 68% of hold-out chemicals have predicted dose-response curve lower bound that at some point exceeds 0. The five highest priority chemicals using the max lower bound method would be Basic Blue 7, Ergocalciferol, Toremifene citrate, 4-(2-Phenylpropan-2-yl)-N-[4-(2-phenylpropan-2-yl)phenyl]aniline, and 2,4-Bis(1-methyl-1-phenylethyl)phenol. As discussed previously, Toremifene citrate is a close relative of Clomiphene citrate (1:1) (activity shown in Figure 17) and a selective estrogen receptor modulator used to treat ovulatory dysfunction in women trying to become pregnant. Basic Blue 7 is labeled as corrosive, an irritant, acutely toxic, and an environmental hazard. The closest neighbors to Basic Blue 7 in the training set are other colorants, Gentian Violet and Malachite green, that have known toxic effects (Docampo and Moreno (1990), Srivastava, Sinha and Roy (2004)).

## 5. Conclusion.

We have focused on the utility of distance learning for designing future chemical test sets, but these pairwise distance matrices could be used in place of Euclidean distance in any distance-based statistical analysis. This would include distance-based clustering of chemicals as well as kernel and Gaussian process-based models. As a specific example the authors believe these activity-relevant distances have the potential to improve main effects estimates in mixture models for human health outcomes. It is likely that incorporating knowledge about similarity in activity-relevant space (e.g., by using the toxicity-relevant pairwise distance matrix to inform a group penalized regression model) would provide stabilization for main effects and, in turn, allow for better estimation of the interaction effects.

The designation of active vs. inactive in the BS<sup>3</sup>FA model is based on a posterior summary of the predicted dose-response profiles; there is no direct incorporation of the concept of a chemical being inactive in the model itself. It may be desirable, particularly when considering assays having very few chemicals presenting with any activity, to probabilistically model inactivity. For example, the dose-response profile could be modeled as a mixture between the zero-vector and the BS<sup>3</sup>FA factor model with a learned weight on the zero vector corresponding to the probability of inactivity.

The BS<sup>3</sup>FA model deals with the structured decomposition of a single assay and a single feature data set. In reality, any sort of model hoping to extend to human health outcomes will need to utilize information from multiple sources. In the ToxCast data set, there is not just one dose-response curve per chemical. There are many assay endpoints of potential relevance to human toxicity. Furthermore, there are potentially many useful chemical feature descriptors (we used Mold2, but others include MACCS keys, Daylight Fingerprints, or Morgan Fingerprints, to name a few). Future work will link the ideas in BS<sup>3</sup>FA to those in Wilson, Reif and Reich (2014) to allow for a more direct active/inactive assignment and to hierarchically describe variability across multiple assays. Extending even further, the holy grail of toxicity modeling would be an explicit linking of multiple assay endpoints to human health data such that human health outcomes could be predicted from chemical structure alone.

## Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

## Acknowledgments.

The authors would like to thank Evan Poworoznek and Bora Jin for helpful comments.

## Funding.


This work was partially supported by the National Institute of Environmental Health Sciences of the United States National Institutes of Health (grants 1R01ES028804-01 and 5R01ES027498-02, and intramural funds) and the Department of Energy Computational Science Graduate Fellowship (grant DE-FG02-97ER25308). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

## REFERENCES

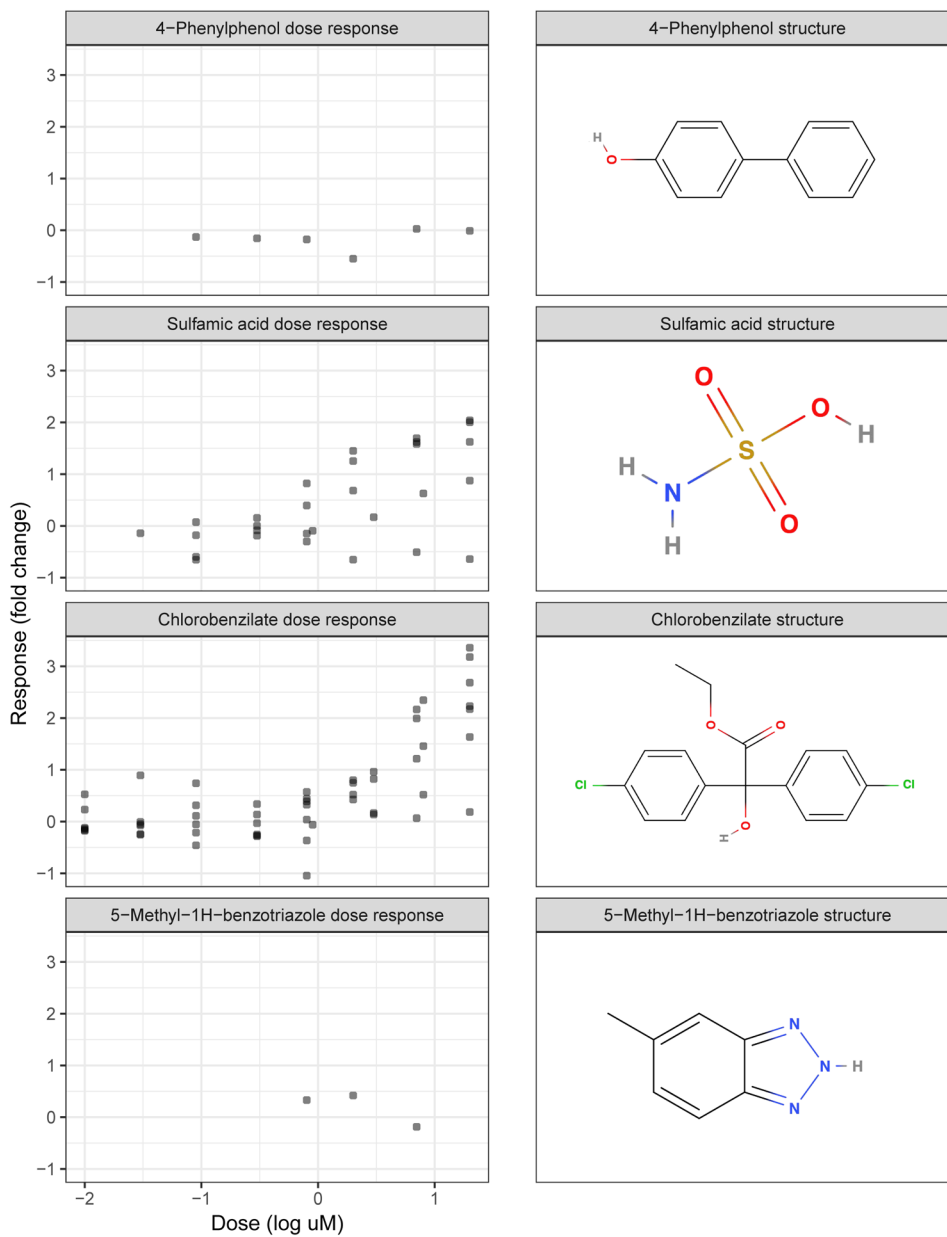
- Barber RF, Reimherr M and Schill T (2017). The function-on-scalar LASSO with applications to longitudinal GWAS. *Electron. J. Stat* 11 1351–1389. MR3635916 10.1214/17-EJS1260
- Bhattacharya A and Dunson DB (2011). Sparse Bayesian infinite factor models. *Biometrika* 98 291–306. MR2806429 10.1093/biomet/asr013 [PubMed: 23049129]
- Canale A and Dunson DB (2013). Nonparametric Bayes modelling of count processes. *Biometrika* 100 801–816. MR3142333 10.1093/biomet/ast037
- Carvalho CM, Polson NG and Scott JG (2010). The horseshoe estimator for sparse signals. *Biometrika* 97 465–480. MR2650751 10.1093/biomet/asq017
- Chen Y, Goldsmith J and Ogden RT (2016). Variable selection in function-on-scalar regression. *Stat* 5 88–101. MR3478799 10.1002/sta4.106 [PubMed: 27429751]
- Dhaliwal LK, Suri V, Gupta KR and Sahdev S (2011). Tamoxifen: An alternative to clomiphene in women with polycystic ovary syndrome. *Journal of Human Reproductive Sciences* 4 76. [PubMed: 22058644]

- Dix DJ, Houck KA, Martin MT, Richard AM, Setzer RW and Kavlock RJ (2007). The ToxCast program for prioritizing toxicity testing of environmental chemicals. *Toxicol. Sci* 95 5–12. 10.1093/toxsci/kfl1103 [PubMed: 16963515]
- Docampo R and Moreno SN (1990). The metabolism and mode of action of gentian violet. *Drug Metab. Rev* 22 161–178. [PubMed: 2272286]
- Durante D (2017). A note on the multiplicative gamma process. *Statist. Probab. Lett* 122 198–204. MR3584158 10.1016/j.spl.2016.11.014
- Fan Z and Reimherr M (2017). High-dimensional adaptive function-on-scalar regression. *Econom. Stat* 1 167–183. MR3669995 10.1016/j.ecosta.2016.08.001
- Hahn PR and Carvalho CM (2015). Decoupling shrinkage and selection in Bayesian linear models: A posterior summary perspective. *J. Amer. Statist. Assoc* 110 435–448. MR3338514 10.1080/01621459.2014.993077
- Hong H, Xie Q, Ge W, Qian F, Fang H, Shi L, Su Z, Perkins R and Tong W (2008). Mold2, molecular descriptors from 2D structures for chemoinformatics and toxicoinformatics. *J. Chem. Inf. Model* 48 1337–1344. [PubMed: 18564836]
- Hong H, Slavov S, Ge W, Qian F, Su Z, Fang H, Cheng Y, Perkins R, Shi L et al. (2012). Mold2 molecular descriptors for QSAR. In *Statistical Modelling of Molecular Descriptors in QSAR/QSPR* 2 65–109.
- Judson RS, Houck KA, Kavlock RJ, Knudsen TB, Martin MT, Mortensen HM, Reif DM, Rotroff DM, Shah I et al. (2009). In vitro screening of environmental chemicals for targeted testing prioritization: The ToxCast project. *Environ. Health Perspect* 118 485–492.
- Kavlock R, Chandler K, Houck K, Hunter S, Judson R, Kleinstreuer N, Knudsen T, Martin M, Padilla S et al. (2012). Update on EPA's ToxCast program: Providing high throughput decision support tools for chemical risk management. *Chem. Res. Toxicol* 25 1287–1302. [PubMed: 22519603]
- Kliwer SA, Goodwin B and Willson TM (2002). The nuclear pregnane X receptor: A key regulator of xenobiotic metabolism. *Endocr. Rev* 23 687–702. 10.1210/er.2001-0038 [PubMed: 12372848]
- Knowles D and Ghahramani Z (2011). Nonparametric Bayesian sparse factor models with application to gene expression modeling. *Ann. Appl. Stat* 5 1534–1552. MR2849785 10.1214/10-AOAS435
- Kowal DR and Bourgeois DC (2020). Bayesian function-on-scalars regression for high-dimensional data. *J. Comput. Graph. Statist* 29 629–638. MR4153187 10.1080/10618600.2019.1710837
- Li G, Shen H and Huang JZ (2016). Supervised sparse and functional principal component analysis. *J. Comput. Graph. Statist* 25 859–878. MR3533642 10.1080/10618600.2015.1064434
- Liu R, Rallo R, George S, Ji Z, Nair S, Nel AE and Cohen Y (2011). Classification NanoSAR development for cytotoxicity of metal oxide nanoparticles. *Small* 7 1118–1126. [PubMed: 21456088]
- Lock EF, Hoadley KA, Marron JS and NOBEL AB (2013). Joint and individual variation explained (JIVE) for integrated analysis of multiple data types. *Ann. Appl. Stat* 7 523–542. MR3086429 10.1214/12-AOAS597 [PubMed: 23745156]
- Low-Kam C, Telesca D, Ji Z, Zhang H, Xia T, Zink JI and Nel AE (2015). A Bayesian regression tree approach to identify the effect of nanoparticles' properties on toxicity profiles. *Ann. Appl. Stat* 9 383–401. MR3341120 10.1214/14-AOAS797
- Makalic E and Schmidt DF (2016). A simple sampler for the horseshoe estimator. *IEEE Signal Process. Lett* 23 179–182.
- Martin YC, Kofron JL and Traphagen LM (2002). Do structurally similar molecules have similar biological activity? *J. of Med. Chem* 45 4350–4358. [PubMed: 12213076]
- Meng J, Zhang J, Qi Y, Chen Y and Huang Y (2010). Uncovering transcriptional regulatory networks by sparse Bayesian factor model. *EURASIP J. Adv. Signal Process* 2010 3.
- Meyer MJ, Coull BA, Versace F, Cinciripini P and Morris JS (2015). Bayesian function-on-function regression for multilevel functional data. *Biometrics* 71 563–574. MR3402592 10.1111/biom.12299 [PubMed: 25787146]
- Montagna S, Tokdar ST, Neelon B and Dunson DB (2012). Bayesian latent factor regression for functional and longitudinal data. *Biometrics* 68 1064–1073. MR3040013 10.1111/j.1541-0420.2012.01788.x [PubMed: 23005895]

- Moran KR, Dunson D, Wheeler MW and Herring AH (2021). Supplement to “Bayesian joint modeling of chemical structure and dose response curves.” <https://doi.org/10.1214/21-AOAS1461SUPPA>, <https://doi.org/10.1214/21-AOAS1461SUPPB>, <https://doi.org/10.1214/21-AOAS1461SUPPC>
- Nikolova N and Jaworska J (2003). Approaches to measure chemical similarity—a review. *QSAR & Combinatorial Science* 22 1006–1026.
- O’connell MJ and Lock EF (2019). Linked matrix factorization. *Biometrics* 75 582–592. MR3999181 10.1111/biom.13010 [PubMed: 30516272]
- Patel T, Telesca D, Low-Kam C, Ji ZX, Zhang HY, Xia T, Zinc JI and Nel AE (2014). Relating nanoparticle properties to biological outcomes in exposure escalation experiments. *Environmetrics* 25 57–68. MR3233744 10.1002/env.2246 [PubMed: 24764692]
- Pati D, Bhattacharya A, Pillai NS and Dunson D (2014). Posterior contraction in sparse Bayesian factor models for massive covariance matrices. *Ann. Statist* 42 1102–1130. MR3210997 10.1214/14-AOS1215
- Ray P, Zheng L, Lucas J and Carin L (2014). Bayesian joint analysis of heterogeneous genomics data. *Bioinformatics* 30 1370–1376. [PubMed: 24489367]
- Seyedoshohadaei F, Zandvakily F and Shahgeibi S (2012). Comparison of the effectiveness of clomiphene citrate, tamoxifen and letrozole in ovulation induction in infertility due to isolated unovulation. *Iran. J. Reprod. Med* 10 531–536. [PubMed: 25246922]
- Srivastava S, Sinha R and Roy D (2004). Toxicological effects of malachite green. *Aquat. Toxicol* 66 319–329. [PubMed: 15129773]
- Weininger D (1988). SMILES, a chemical language and information system. I. Introduction to methodology and encoding rules. *J. Chem. Inf. Comput. Sci* 28 31–36.
- Wheeler MW (2019). Bayesian additive adaptive basis tensor product models for modeling high dimensional surfaces: An application to high-throughput toxicity testing. *Biometrics* 75 193–201. MR3953720 [PubMed: 30081432]
- Wilson A, Reif DM and Reich BJ (2014). Hierarchical dose-response modeling for high-throughput toxicity screening of environmental chemicals. *Biometrics* 70 237–246. MR3251684 10.1111/biom.12114 [PubMed: 24397816]
- Yoshida R and West M (2010). Bayesian learning in sparse graphical factor models via variational mean-field annealing. *J. Mach. Learn. Res* 11 1771–1798. MR2653356 [PubMed: 20890391]




**Fig. 1.** Pairwise correlation between each of the 777 molecular descriptors in ToxCast for the chemicals profiled.

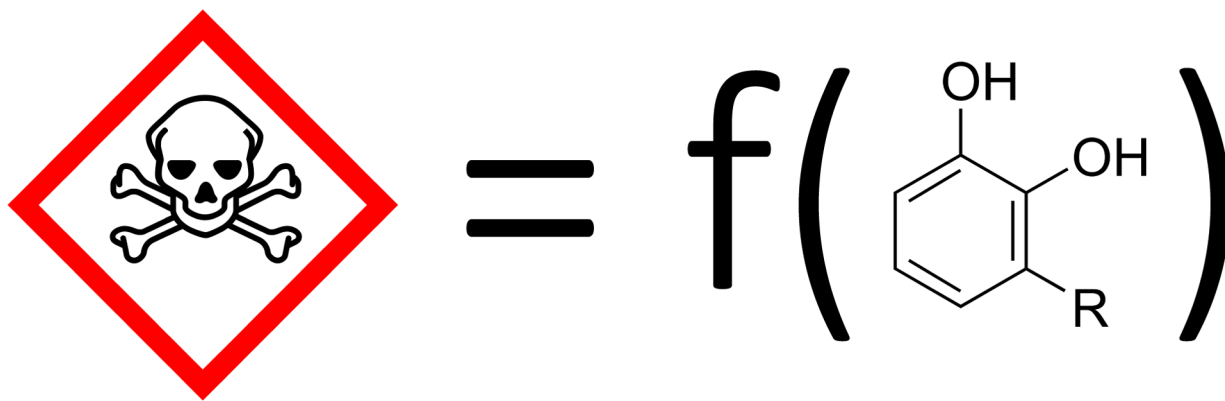


**Fig. 2.** Left: Dose response data for example chemicals from the ToxCast ATG PXR assay (i.e.,  $y_i$ ). Right: 2D chemical structure diagrams for example chemicals (converted from SMILES into  $x_i$  using the Mold2 software).






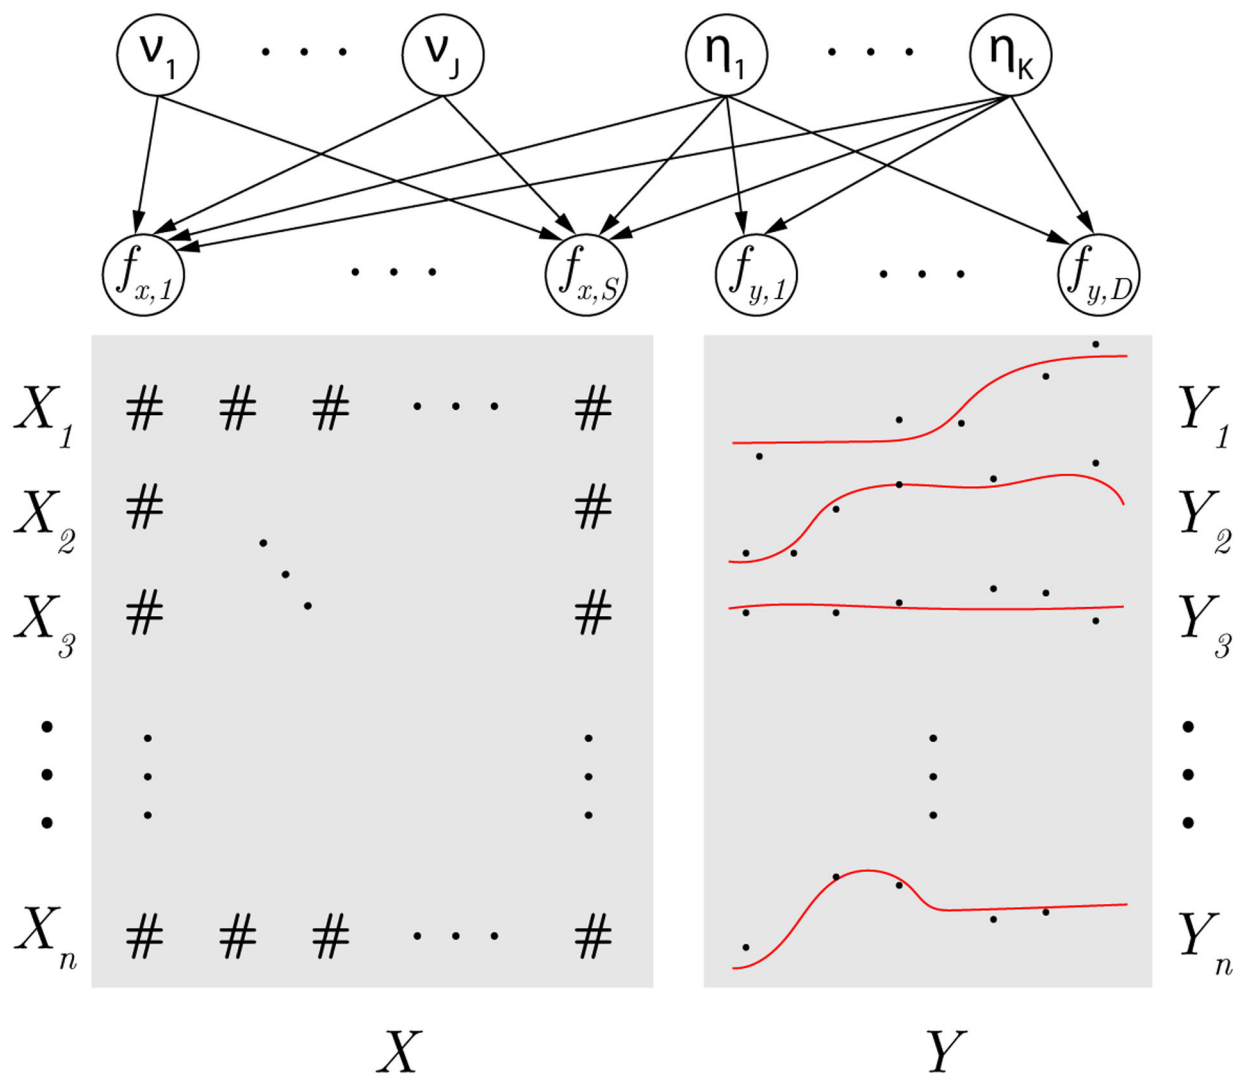
**Fig. 3.** Numeric values for select Mold2 traits of Bisphenol A (BPA). The chemical formula for BPA is  $C_{15}H_{16}O_2$ , and its SMILES descriptor is CC(C)(C1=CC=C(C=C1)O)C2=CC=C(C=C2)O.



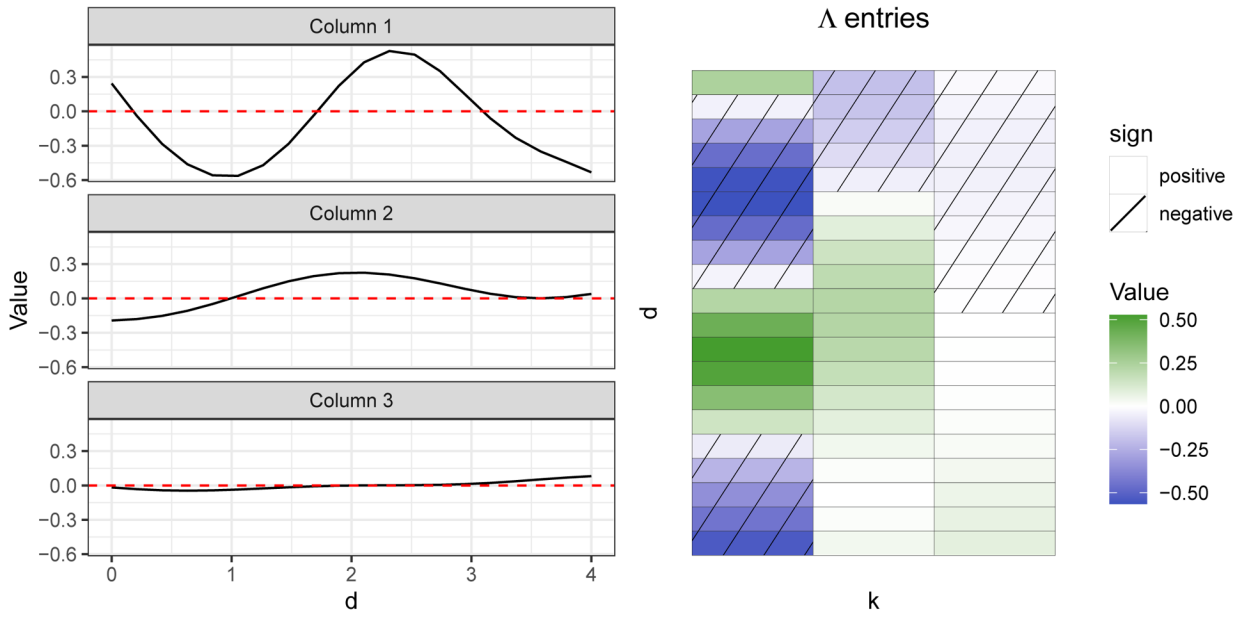
**Fig. 4.** Quantitative Structure-Activity Relationship (QSAR) models predict toxicity as a function of chemical structure.



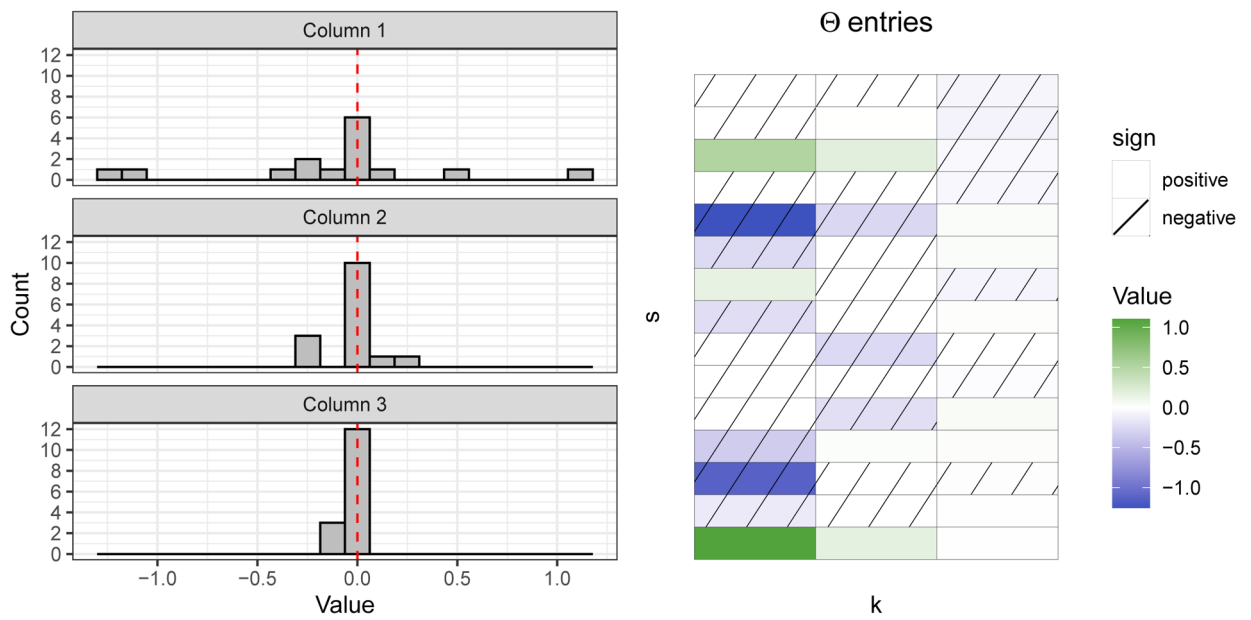
**Fig. 5.** Relationship between chemical “structure distance” and chemical “activity distance” when PCA and functional PCA (FPCA) are performed independently on the Mold2 chemical structures and the ToxCast dose-response curves. Each point on the graph shows the Euclidean distance between the (functional) principal component scores accounting for 95% of the variability in the data for one pair of chemicals on the (y-) x-axis.

**Fig. 6.**

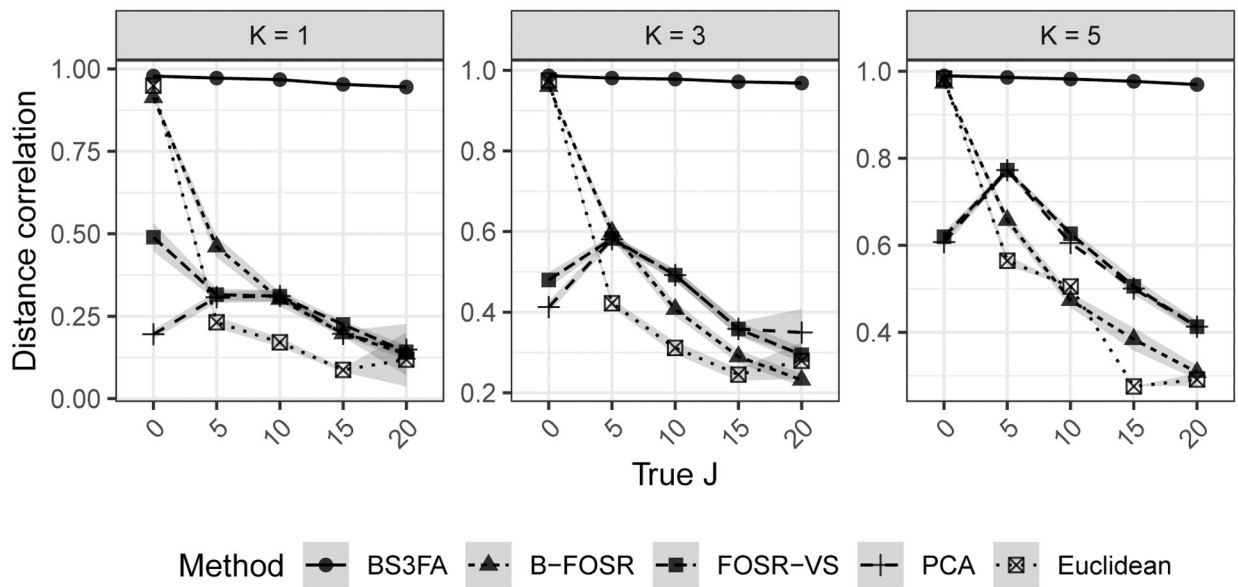
Visual representation of the BS<sup>3</sup>FA model. Entries in  $X$  are the chemical feature descriptors from Mold2, while entries in  $Y$  are noisy realizations of an underlying smooth dose-response curve. Observations  $f_{x,1}, \dots, f_{x,S}$  are the true mean of the chemical features and  $f_{y,1}, \dots, f_{y,D}$  are the true mean of the dose-response values. Latent variables  $\eta_1, \dots, \eta_K$ , the underlying toxicity-relevant factors, are shared, and  $v_1, \dots, v_K$ , the underlying toxicity-irrelevant factors, are specific to the chemical features. Arrows denote probabilistic dependency.



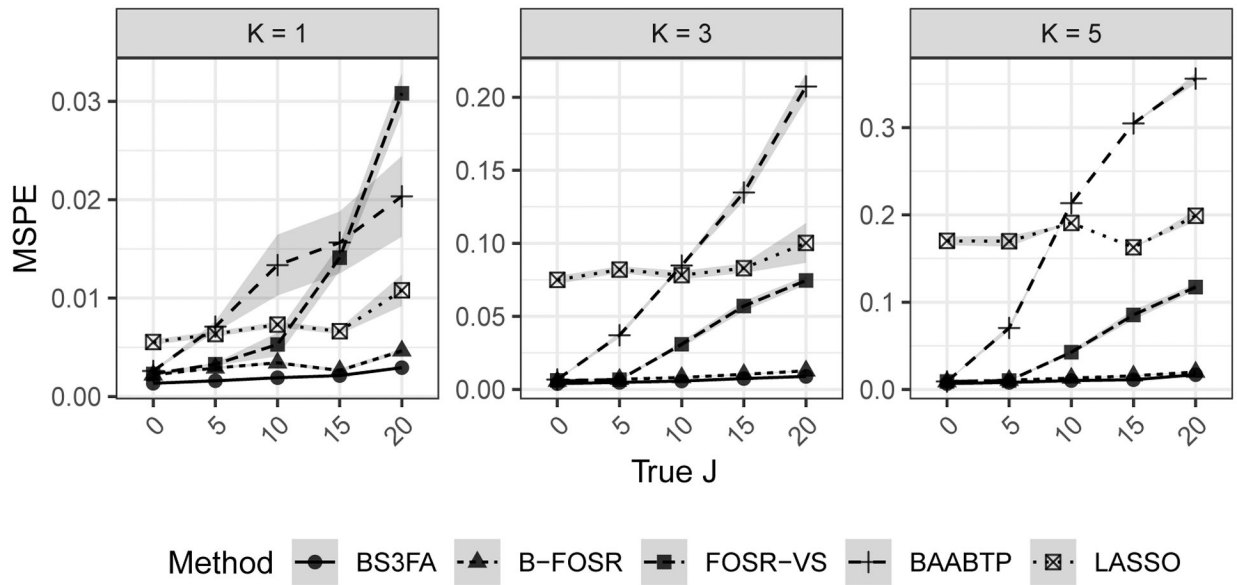
**Fig. 7.** Visualization of the smoothness and the shrinkage of columns of an example  $D = 20$  by  $K = 3$  loadings matrix  $\Lambda$ . The effect is an automatic truncation of the number of factors in the model and a learned latent dimension. In this figure, columns of  $\Lambda$  were sampled using the kernel in equation (6) with  $\ell = 1$  and  $\alpha_k^2$  being  $1, \frac{1}{25}, \frac{1}{400}$  for  $k = 1, 2, 3$ , and the indices  $d$  corresponding to the length-20 vector from 0 to 4, inclusive. Note that as  $\alpha_k^2$  decreases, the functions tend to flatten.



**Fig. 8.** Visualization of elementwise sparsity and the shrinkage of columns of an example  $S = 15$  by  $K = 3$  loadings matrix  $\Theta$ . The effect is an automatic truncation of the number of factors in the model and a learned latent dimension. In this figure, columns of  $\Theta$  had column-specific variance  $\tau_k^{-1}$  being  $1, \frac{1}{25}, \frac{1}{400}$  for  $k = 1, 2, 3$ .



**Fig. 9.** Correlation between entries in the true pairwise distance matrix (i.e., the Euclidean distance between true latent factors  $\eta$ ) and the predicted pairwise distance matrix for holdout chemicals. Each subplot shows the result of 100 simulations per J across methods for a given true shared subspace dimension K.



**Fig. 10.** Mean squared predictive error (MSPE) for the hold-out chemicals' dose-response mean functions. Each subplot shows the result of 100 simulations per  $J$  across methods for a given true shared subspace dimension  $K$ .


Author Manuscript

Author Manuscript


Author Manuscript

Author Manuscript



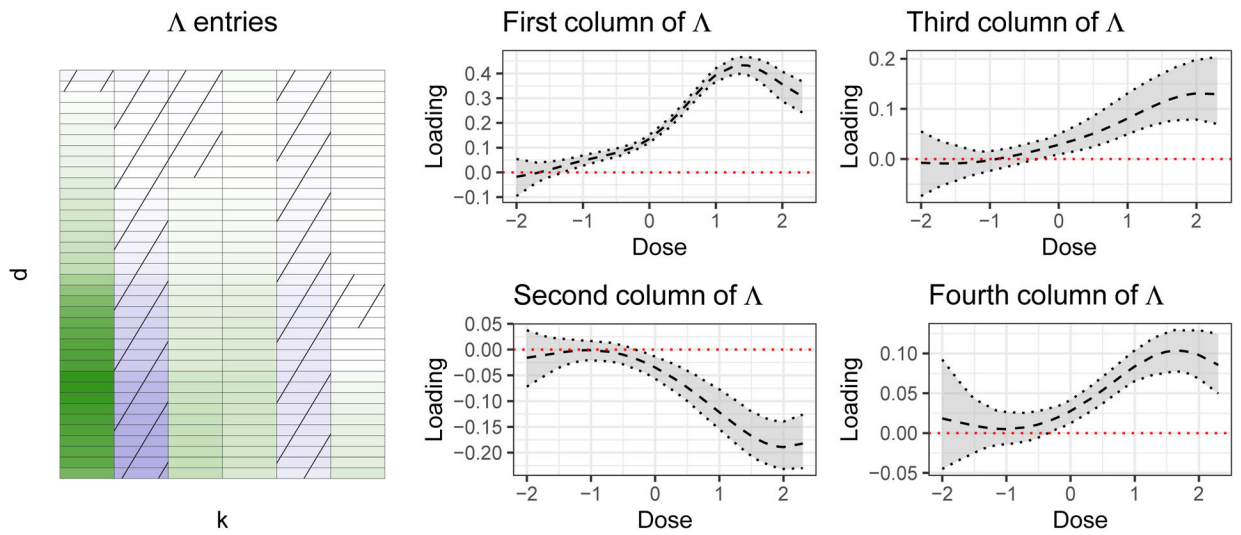


**Fig. 11.** Missingness by dose for each chemical in the analysis. Each row is corresponds to a chemical and each column to a dose value (sorted in ascending order). Note that very few chemicals have observations below  $-2 \log \mu\text{M}$ .




**Fig. 12.**

All response values at each dose in the data set. The solid line shows the model-predicted global mean  $\mu^Y$  and the dashed lines show the model-predicted 95% simultaneous credible band about  $\mu^Y$ .

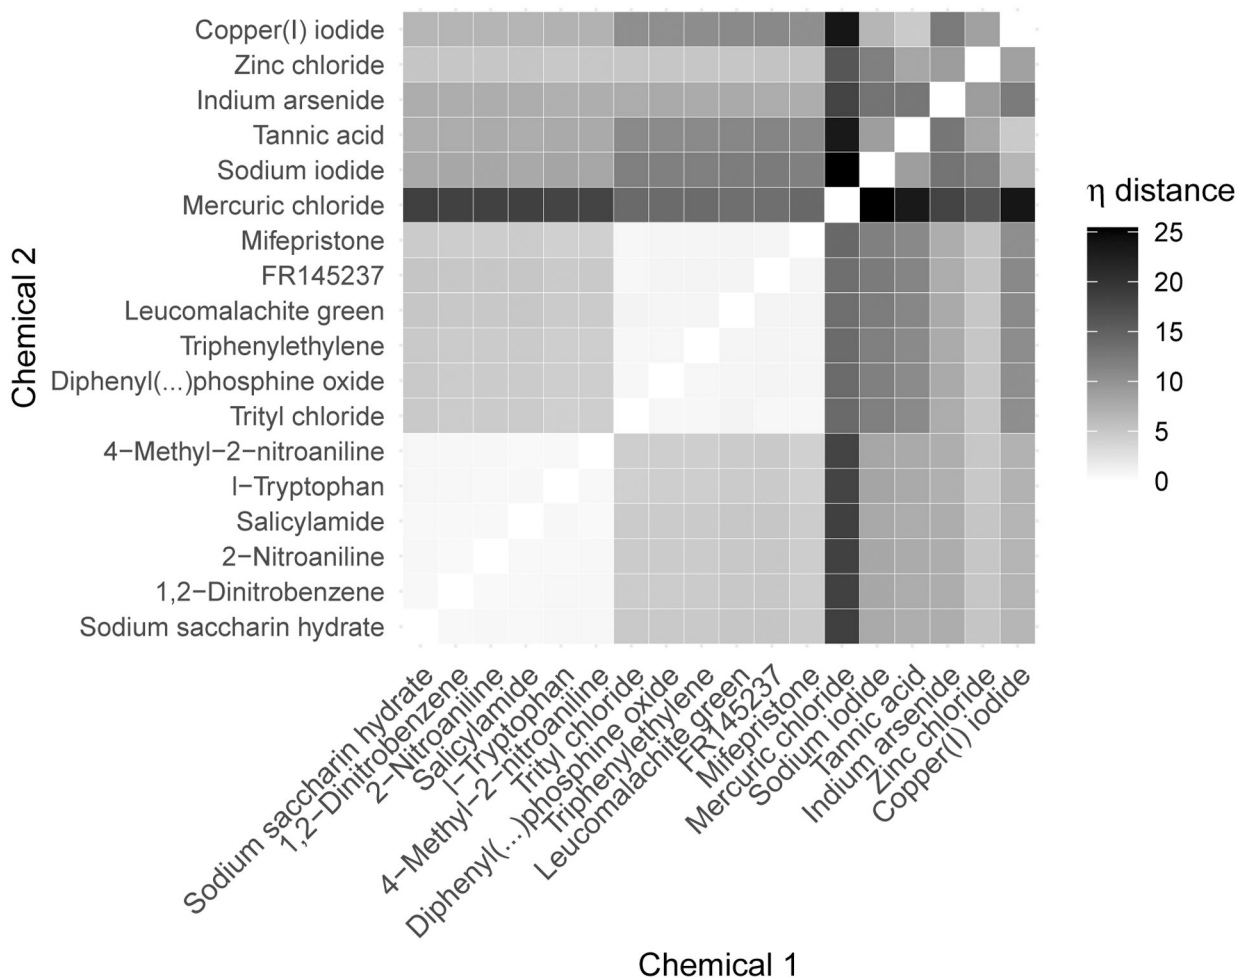
**Fig. 13.**

Left: First six columns of the predicted mean of  $\Lambda$ . Each row represents a unique dose value.


Middle and right: Mean and 95% simultaneous bands for the first four columns of  $\Lambda$ . The dose values are given on the x-axis.



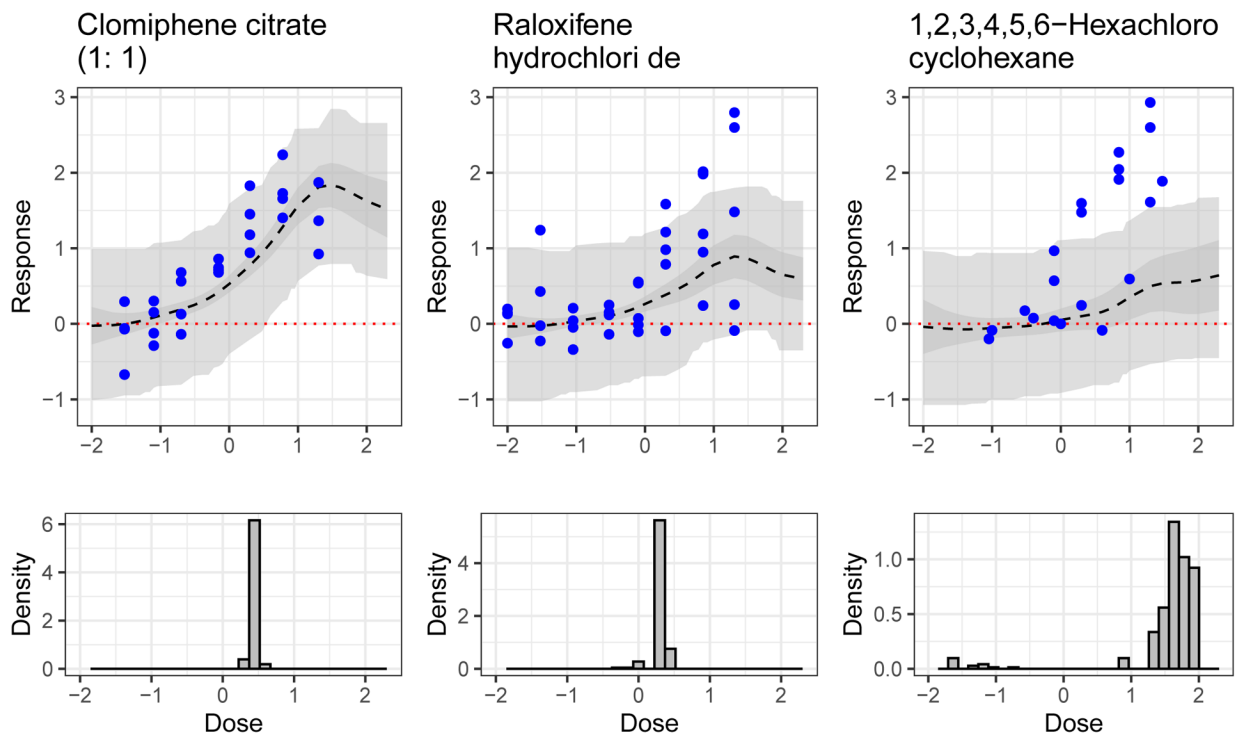
**FIG. 14.** Histogram showing entries of the predicted mean of  $\Theta$ . The bulk of the estimated entries are close to 0 due to the shrinkage effect of the horseshoe prior.



**Fig. 15.** Expected distance in  $\eta$  space for a select set of training chemicals.

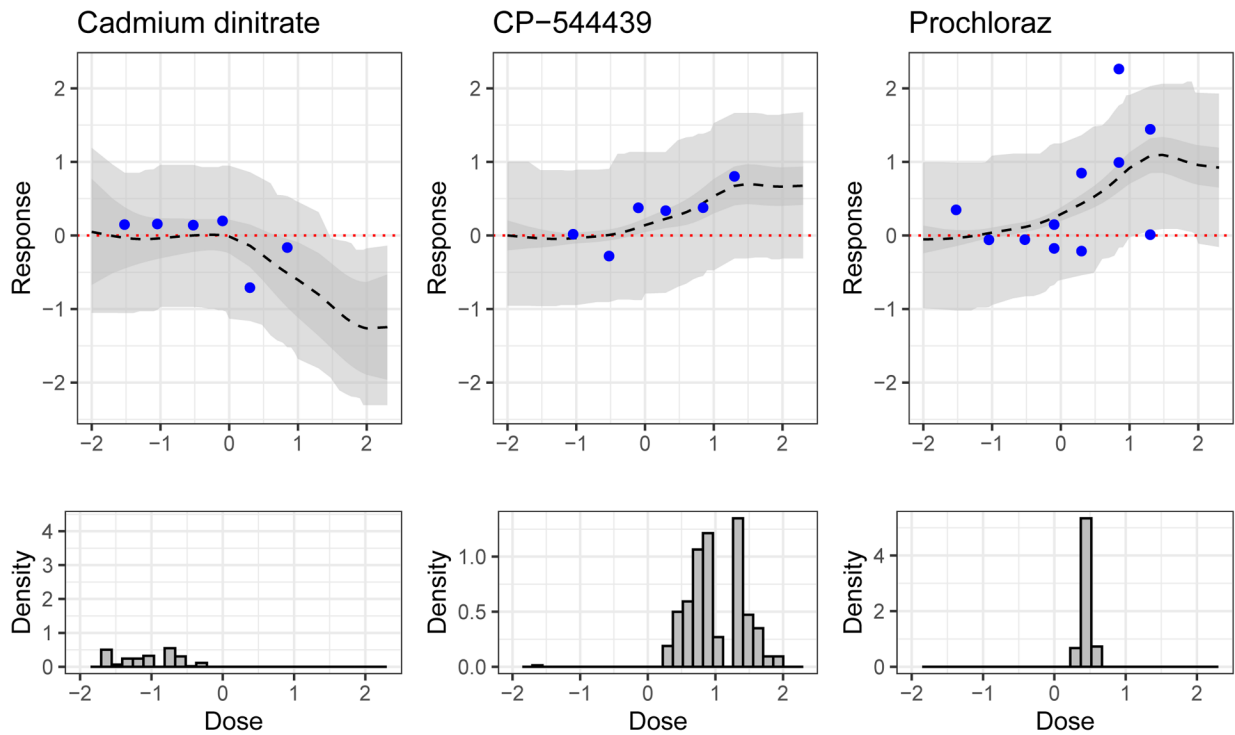


**Fig. 16.** From left to right: Tetracosafuorotetradecahydrophenanthrene, Strychnine hemisulphate salt and Cadmium dinitrate.



**Fig. 17.**

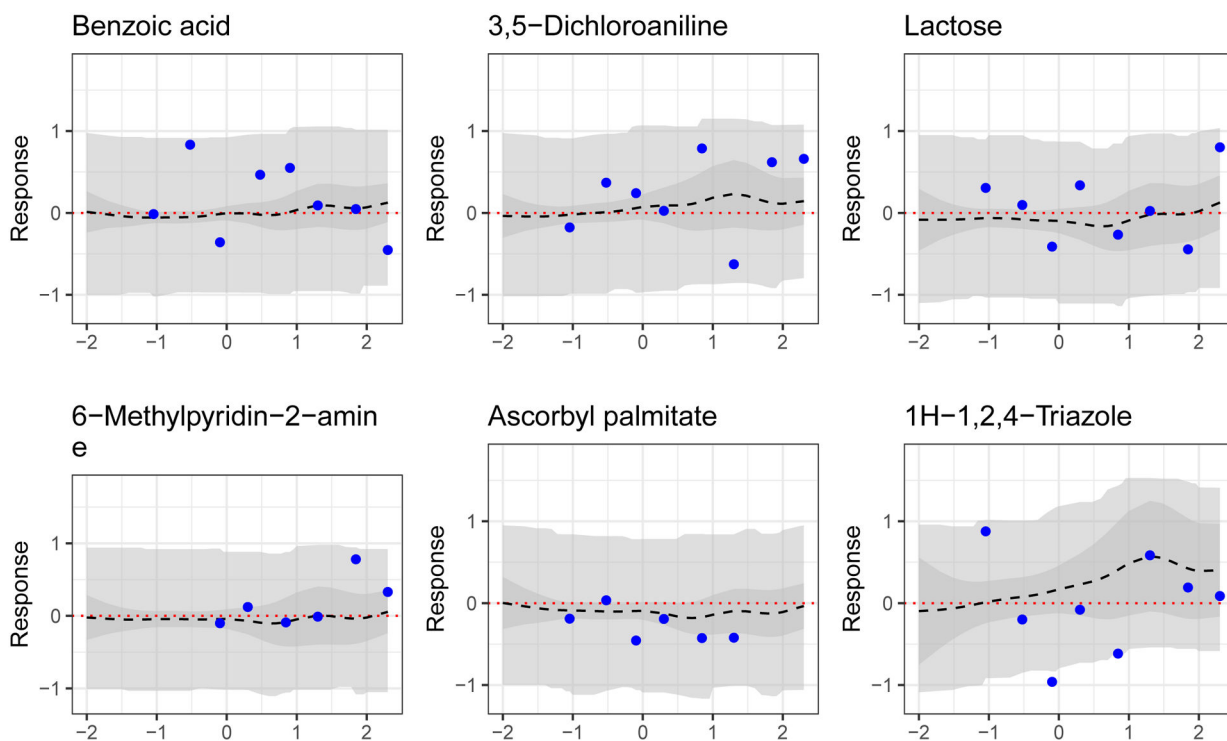
Results for heavily tested hold-out chemicals predicted by the model to be activating. MSEs from left to right are 0.21, 0.49 and 1.44. Top: Predicted average dose-response curve (dashed black line), 95% simultaneous band for expected dose-response curve (darker grey ribbon) and 95% simultaneous band for observed data (lighter grey ribbon). Data (held out in training) are solid points. Bottom: Posterior samples of the AC50 value, that is, the dose at which the dose-response curve is at half of its maximal value.



**Fig. 18.**

Results for hold-out chemicals predicted by the model to be active (activity decreasing in the case of Cadmium dinitrate). MSEs from left to right are 0.10, 0.03 and 0.40. Top: Predicted average dose-response curve (dashed black line), 95% simultaneous band for expected dose-response curve (darker grey ribbon) and 95% simultaneous band for observed data (lighter grey ribbon). Data (held out in training) are solid points. Bottom: Posterior samples of the AC50 value, that is, the dose at which the dose-response curve is at half of its maximal value.





**Fig. 19.**

Results for select hold-out chemicals predicted by the model to be inactive. MSEs from left to right, top to bottom, are 0.22, 0.23, 0.14, 0.13, 0.06 and 0.42. Shown are predicted average dose-response curve (dashed black line), 95% simultaneous band for expected dose-response curve (darker grey ribbon) and 95% simultaneous band for observed data (lighter grey ribbon). Data (held out in training) are solid points.

**Table 1**

True positive rate (TPR), false positive rate (FPR) and false discovery rate (FDR) for the proposed method of assessing whether a chemical is active. A perfect classifier has a TPR of 1 and an FPR/FDR of 0. Shown is the mean (SD) across simulations

		$J = 0$	$J = 5$	$J = 10$	$J = 15$	$J = 20$
TPR	$K = 1$	0.72 (0.08)	0.66 (0.08)	0.64 (0.07)	0.58 (0.06)	0.52 (0.09)
	$K = 3$	0.98 (0.03)	0.97 (0.03)	0.97 (0.03)	0.95 (0.04)	0.93 (0.05)
	$K = 5$	1.00 (0.01)	1.00 (0.01)	1.00 (0.01)	0.99 (0.01)	0.98 (0.02)
FPR	$K = 1$	0.00 (0.00)	0.00 (0.00)	0.00 (0.01)	0.00 (0.00)	0.00 (0.00)
	$K = 3$	0.00 (0.01)	0.00 (0.01)	0.00 (0.01)	0.00 (0.01)	0.00 (0.01)
	$K = 5$	0.00 (0.01)	0.00 (0.01)	0.00 (0.01)	0.00 (0.01)	0.00 (0.01)
FDR	$K = 1$	0.00 (0.01)	0.00 (0.01)	0.00 (0.01)	0.00 (0.00)	0.00 (0.01)
	$K = 3$	0.00 (0.01)	0.00 (0.01)	0.00 (0.01)	0.00 (0.01)	0.00 (0.01)
	$K = 5$	0.00 (0.01)	0.00 (0.01)	0.00 (0.01)	0.00 (0.01)	0.00 (0.01)