

# ALOHA: Aggregated local extrema splines for high-throughput dose–response analysis

Sarah E. Davidson<sup>a,\*</sup>, Matthew W. Wheeler<sup>b,\*</sup>, Scott S. Auerbach<sup>c</sup>, Siva Sivaganesan<sup>d</sup>, Mario Medvedovic<sup>e</sup>

<sup>a</sup> Department of Environmental Health Sciences Division of Biostatistics and Bioinformatics, University of Cincinnati, Cincinnati, OH, United States

<sup>b</sup> Division of Intramural Research (DIR) Biostatistics & Computational Biology Branch (BCBB), National Institute of Environmental Health Sciences (NIEHS), Durham, NC, United States

<sup>c</sup> National Toxicology Program Division (DNT) Biomolecular Screening Branch (BSB), National Institute of Environmental Health Sciences (NIEHS), Durham, NC, United States

<sup>d</sup> Department of Mathematical Sciences Division of Statistics and Data Science, University of Cincinnati, Cincinnati, OH, United States

<sup>e</sup> Department of Environmental Health Sciences Division of Biostatistics and Bioinformatics, University of Cincinnati, Cincinnati, OH, United States

## ARTICLE INFO

### Keywords:

Biological pathways  
Genomic benchmark dose  
High throughput data  
Bayesian clustering

## ABSTRACT

Computational methods for genomic dose–response integrate dose–response modeling with bioinformatics tools to evaluate changes in molecular and cellular functions related to pathogenic processes. These methods use parametric models to describe each gene's dose–response, but such models may not adequately capture expression changes. Additionally, current approaches do not consider gene co-expression networks. When assessing co-expression networks, one typically does not consider the dose–response relationship, resulting in 'co-regulated' gene sets containing genes having different dose–response patterns. To avoid these limitations, we develop an analysis pipeline called Aggregated Local Extrema Splines for High-throughput Analysis (ALOHA), which computes individual genomic dose–response functions using a flexible class Bayesian shape constrained splines and clusters gene co-regulation based upon these fits. Using splines, we reduce information loss due to parametric lack-of-fit issues, and because we cluster on dose–response relationships, we better identify co-regulation clusters for genes that have co-expressed dose–response patterns from chemical exposure. The clustered pathways can then be used to estimate a dose associated with a pre-specified biological response, i.e., the benchmark dose (BMD), and approximate a point of departure dose corresponding to minimal adverse response in the whole tissue/organism. We compare our approach to current parametric methods and our biologically enriched gene sets to cluster on normalized expression data. Using this methodology, we can more effectively extract the underlying structure leading to more cohesive estimates of gene set potency.

## Introduction

Genomic dose–response (GDR) modeling evaluates the relationship between chemical exposure and transcriptional responses to identify a level of minimal biological effect, called a point-of-departure (POD).

Traditionally, PODs are estimated using *in vivo* toxicology responses. It has been proposed when minimal *in vivo* data exists, GDR data can be used in an interim risk assessment to estimate a genomic POD [1–4]. Acceptance of this approach relies on mitigating concerns about the qualitative differences between genomic and traditional PODs. A

**Abbreviations:** ALOHA, Aggregated Local extrema splines for High-throughput Analysis; AD, All Doses; AIC, Akaike Information Criterion; BAP, Benzo-[ $\alpha$ ]-pyrene; BLX, Bayesian Local Extrema; BMD, Benchmark Dose; BMDL, lower BMD uncertainty bound; BMDU, upper BMD uncertainty bound; BMR, Benchmark Response; BPF, Best Parametric Fit; CDR, Cytotoxic Doses Removed; CGN, Cancer Gene Neighborhoods; CLEAN, Cluster Enrichment Analysis; CPM, Count per Million; CPZ, Chlorpromazine; CSIMM, context-specific mixture models; DR, dose-response; DPM, Dirichlet Process Mixture Model; ECDF, Empirical Cumulative Distribution Function; FFDR, Fisher's exact test Benjamini-Hochberg False Discovery Rate adjusted p-value; GDR, genomic dose-response; KCl, Potassium chloride; LODOCV, leave-one-dose-out cross-validation; MCMC, Markov Chain Monte Carlo; MSigDB, Molecular Signature Database; MSPE, Mean squared predicted error; MRC, Minimum Read Count; POD, point-of-departure.

\* Corresponding authors.

E-mail addresses: [davidsss@mail.uc.edu](mailto:davidsss@mail.uc.edu) (S.E. Davidson), [matt.wheeler@nih.gov](mailto:matt.wheeler@nih.gov) (M.W. Wheeler).

<https://doi.org/10.1016/j.comtox.2021.100196>

Available online 13 October 2021

2468-1113/© 2021 Elsevier B.V. All rights reserved.

suitable estimation of transcriptional PODs accounts for differences in biological relevance between perturbed pathways and incorporates mechanistic information [2,5].

To date, several GDR modeling approaches have been developed with a variety of methodological features, see S1 Table 1. Except for BMDEExpress [6–8], which has documented guidelines in National Toxicology Program (NTP) et al. [9], none of the other approaches have established an approach for determining a genomic POD linked to gene- and pathway-level biological alterations. Consequently, BMDEExpress [6–8] is still the most widely utilized approach for GDR modeling and will be the focus of our comparison. The platform uses a multi-step framework of parametric modeling and bioinformatics tools to characterize dose–response (DR) patterns and evaluate alterations in cellular function [6–9]. It addresses common challenges by using genomic data in risk assessment, but some of the assumptions made may not be appropriate given the observed data. In cases where cytotoxicity occurs in high-doses, a non-monotonic (i.e., curve with umbrella orderings) is possible [10]. For such shapes, parametric models are typically inadequate, because most of the models assume monotone changes in expression [9,11–13], which may lead to removing genes with biological relevance due to poor model fit. Another limitation of BMDEExpress is the enrichment analysis used to detect altered cellular functions and obtain genomic PODs. The enrichment analysis matches genes remaining after model fitting to a list of curated gene sets, such as the Hallmark gene sets [14–16], determined a priori. The remaining genes' biological relevance is evaluated by their association to these sets. However, relying on curated gene sets disregards the underlying correlation structure of gene responses. These correlations may clarify understanding of mechanisms, such as co-regulation, related to developing a toxicological event.

Allocco et al. [17] demonstrated grouped genes with highly correlated expressions are likely to indicate co-regulation by a common transcription factor. Clustering-based upon similar dose–response patterns may identify co-regulation. Studies have shown Bayesian clustering methods employing Dirichlet Process mixture (DPM) models are preferred for identifying clusters of co-regulated genes, see Dahl [18], Fraley and Raftery [19], Medvedovic and Sivaganesan [20], Medvedovic et al. [21], Yeung et al. [22] among others. These methods do not always consider relationships among samples that may influence co-regulation patterns. DPM clustering methods that incorporate the relationship among samples using latent variables are shown to improve clustering, see Bar-Joseph et al. [23], Kirk et al. [24], Liu et al. [25], Lock and Dunson [26], Luan and Li [27,28], McDowell et al. [29], Park and Kyung [30], Savage et al. [31] and references therein. However, there is not currently an approach that accounts for the underlying correlation structure of co-expressed genes with similar DR patterns.

Despite the advantage of a latent DPM model that simultaneously includes DR modeling and gene clustering, it is well known these approaches are computationally intensive. Our method, Aggregated Local Extrema Splines for High-throughput dose–response Analysis (ALOHA),

**Table 1**

Three approaches for clustering genes to identify co-expression clusters with similar DR patterns.

Strategy (Abbreviation)	Data	CSIMM Contexts
Original DR Data (CSIMM)	Normalized expression data with all replicate samples	Dose groups $d \in \{1, \dots, D\}$ $M = D$
BLX Estimated DR Data (ALOHA)	Average estimated expression for each dose group using individual BLX fits	Dose groups $d \in \{1, \dots, D\}$ $M = D$
BPF Estimated DR Data (BMDEExpress)	Average estimated expression for each dose group using individual best parametric fits (BPF)	Dose groups $d \in \{1, \dots, D\}$ $M = D$

first fits DR data, then uses a simpler DPM to cluster genes sharing similar DR patterns. From the DPM analysis, the method uses individual clusters to define PODs for co-regulated gene networks. The first stage of our analysis pipeline fits a flexible DR model using Bayesian local extremum splines [32], which is a shape-constrained non-parametric spline approach providing flexibility in fitting monotone, umbrella, and 'J' shaped responses. The second stage of our analysis inputs these fits into the context-specific infinite mixture model (CSIMM) developed by Liu et al. [25], and clusters genes with similar expression patterns. By utilizing gene fits, we incorporate GDR data's dose-dependent nature for clustering while removing the scatter typically observed in such experiments. Biologically enriched gene sets are evaluated to obtain a POD possibly related to adverse biological responses.

ALOHA consists of two parts (i) shape-constrained non-parametric modeling on individual genes discussed in Section 2.1 and (ii) clustering genes with similar DR patterns covered in Section 2.2. In Section 2.3, we outline the application of our method to genomic dose–response data from HepaRG cells exposed to various compounds [33] and compare the results with the current GDR modeling and clustering approaches, respectively.

## Methods

### Dose–Response Modeling

We observe  $n$  genes across  $D$  dose groups, with each dose group having  $K$  replicates such that  $K \geq 1$ . For simplicity of exposition,  $K$  is constant across each dose group but could change at each dose level. Thus, we have  $p = D \times K$  total observations for each gene. Let  $Y$  be the  $n \times p$  matrix of these observations, such that the  $i^{\text{th}}$  gene is represented by the row-vector  $y_i = (y_{(i,1)}, \dots, y_{(i,p)})$ . The expression of gene  $i$  over the experimental doses can be represented as:

$$y_i = f_i(x) + \epsilon_i \quad (1)$$

where  $x = (x_1, \dots, x_p)$  is a vector of doses,  $f_i$  is the DR function, and  $\epsilon_i = (\epsilon_{(i,1)}, \dots, \epsilon_{(i,p)})$  is a vector of measurement errors. Here, we assume  $\epsilon_i \sim N(0, \sigma^2 I)$ , where  $I$  is a  $p \times p$  identity matrix and  $\sigma^2$  is some unknown variance. We use Bayesian shape-constrained non-parametric models to estimate the dose–response curve ( $f_i(x)$ ) for all  $n$  genes.

### Shape-constrained Non-parametric Modeling

For shape-constrained non-parametric modeling, we use a variant of Bayesian local extrema (BLX) splines [32]. This approach places restrictions on the shape of the model space allowing functions with no more than  $H$  extrema (i.e., change-points or local minima/maxima). We define the BLX spline to model functions having at most two extrema. That is, they may increase and then decrease (i.e., have an umbrella ordering), decrease and then increase (i.e., have a 'J' shaped ordering), or increase then decrease and increase again (i.e., have a tilde – ' $\sim$ ' – shaped ordering). In practice, we do not consider the ' $\sim$ ' shape plausible, but it is required to consider 'J' and umbrella orderings simultaneously. When a ' $\sim$ ' shape is observed, it typically represents no dose–response relationship, and as a result, the POD for the given gene is greater than the maximum tested doses.

For gene  $i$ , we estimate the dose–response,  $f_i$ , with a BLX spline defined over the fixed knot set  $T = \{0, 0.5, 1\}$ , where knots are chosen to represent the background, midpoint and maximum dose. The approach is different from Wheeler et al. [32] who consider a prior over a set of knots dense in  $[0, 1]$ . For our application, there are not enough dose groups to use this prior, and the fixed set  $T$  is adequate. With  $H = 2$ , the BLX spline model for  $f_i$  is:

$$f_i(x) = \beta_0 + \sum_{k=1}^4 \beta_k B_k^*(x), \quad (2)$$

where  $\beta_0$  is the intercept,  $B^*(x)$  is a BLX spline basis function of degree 2,

and  $\beta_k$  for  $k = \{1, \dots, 4\}$  are spline coefficients. For more information on construction, we refer the reader to Wheeler et al. [32].

We assume dose–response relationships are initially increasing by setting the fixed integer of the BLX spline basis to one ( $M = 1$ ). The function models monotonic increasing, monotonic decreasing, J-shaped, umbrella shaped, or ‘~’ shaped curves (see S1 Appendix Section 3.1 and S1 Figs. 2–7 therein). All model parameters are estimated as in [32].

Though the dose–response of genes may differ, unlike parametric modeling, the BLX spline specification is the same for each gene. This model similarity facilitates clustering on the estimated parameters (i.e., one can cluster the dose–response curve on the BLX fitted parameters). The BLX approach is different from parametric models, where the parameters all have different interpretations making clustering between models difficult.

### Parametric Modeling

For the parametric modeling approach, we estimate  $f_i$  using all available models, excluding the Polynomial 3 and 4, from BMDEExpress (version 2.3) [6–8]. We refer the reader to Auerbach [6], Phillips et al. [7], Yang et. al. [8] for the model equations. Higher-order polynomials in most cases will fit dose–response data similarly to a quadratic model with fewer degrees of freedom and therefore were excluded from our set. This model suite captures a majority of biological response shapes thought to arise in practice. We use the model with the best fit, determined using the Akaike Information Criterion (AIC) Akaike et al. [34], Akaike [35], Sakamoto et al. [36], for inference and comparison to the BLX fit.

Other than using a quadratic dose–response curve, the parametric model suite cannot represent ‘J’ or umbrella-shaped orderings, which cytotoxic doses can induce. A quadratic model – the only parametric model that can evaluate a ‘J’-shaped response – assumes the response is approximately parabolic forcing symmetry. Symmetry around the extrema occurs infrequently. Dose–response relationships that enforce symmetry may be inadequately represented in the low-dose region.

### BLX Spline and Parametric Model Comparison

Spline models may overfit. To investigate this possibility for BLX performance, we utilize leave-one-dose-out cross-validation (LODOCV). Here, we iteratively leave out each non-control dose group and fit the model (BLX or best parametric model) to the remaining data. The fit’s predicted response is then compared to those from the excluded dose group, and the squared residuals are computed for the held-out data only. The average of these squared residuals across fits is the mean squared predicted error (MSPE) for an individual gene. This procedure is done for each gene probe, and it measures the predictive fit between the BLX and parametric models. Explicitly the MSPE for this LODOCV is,

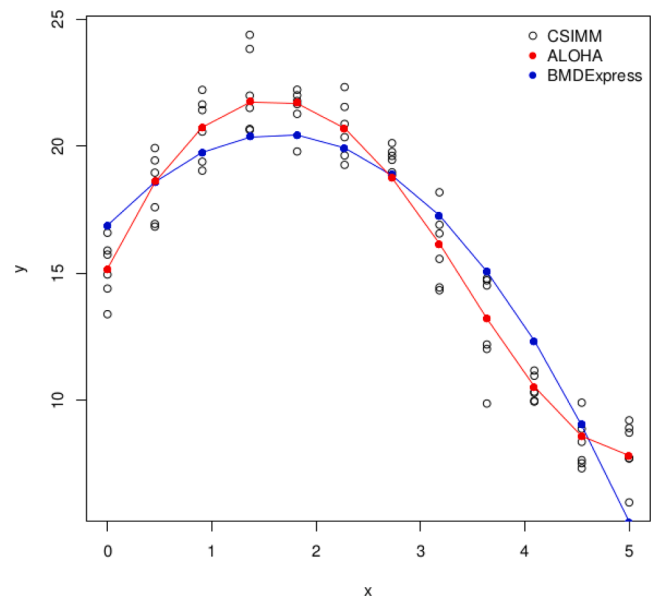
$$MSPE_{(i,m)} = \frac{\sum_{d=1}^D \sum_{j=1}^{n_d} (y_{idj} - f_{(-d)im}(x_j))^2}{\sum_{d=1}^D n_d} \quad (3)$$

where  $i$  denotes the gene,  $m$  denotes the best parametric or BLX spline fit,  $d$  denotes the dose group (excluding the control group),  $n_d$  is the number of samples in  $d$ ,  $j$  denotes the sample,  $x_j$  denotes the dose administered for sample  $j$ , and  $f_{(-d)im}$  is the model fit without dose group  $d$ . Note, in the calculation above, there are no degrees of freedom lost as the held out data are not used in the fit.

### Gene Clustering

We identify co-expression clusters sharing similar DR patterns using the context-specific infinite mixture models (CSIMM) clustering algorithm developed by Liu et al. [25], which allows us to pre-specify sample contexts and evaluate global and local patterns. We modify this approach to model-fit data.

After dose–response modeling, the ALOHA clustering algorithm uses the BLX estimated dose–response data (e.g. red points in Fig. 1) as input



**Fig. 1.** Simulated dose–response data and model fits demonstrating data used for gene clustering approaches. Black points illustrate the observed data values used as input in CSIMM. Red points illustrate the BLX estimated data values used as input in ALOHA. Blue points illustrate the estimated values from a quadratic model fit.  $f(x) = 7\sin(0.3\pi x) + 15 + \epsilon_i$ , where  $\epsilon_i \sim N(0, 1)$ .

for the CSIMM MCMC [25]. We compare the ALOHA clustering algorithm with two other alternatives, including CSIMM and clustering on BMDEExpress model-fit data. CSIMM uses the observed expression data (e.g. black points in Fig. 1) and BMDEExpress uses the estimated dose–response data from the best parametric models (e.g. blue points in Fig. 1). All three algorithms used are given in Table 1. We compare these clustering results to evaluate their ability to capture biologically relevant co-expression clusters with coherent DR patterns and potency estimates (i.e., benchmark dose) for enriched gene sets. For further details on the CSIMM MCMC and the posterior-pairwise probability (PPP) matrix we refer the reader to Section 3.3 in the S1 Appendix and Liu et al. [25].

In addition to the clustering comparisons, we investigate the Hallmark gene sets identified in the pathway enrichment analysis used in BMDEExpress. Enrichment analyses are qualitatively different, and we do not compare the enrichment scores from the two approaches. Our focus was on comparing the identified gene sets, the gene set level potency estimates (i.e., BMD and its interval) obtained from the clustering approach compared to BMDEExpress, and the core genes contributing to the gene set level BMD.

### Application

#### Data

To investigate ALOHA, we use datasets on three compounds tested by Ramaiahgari et al. [33]. Of these compounds, Benzo[a]pyrene (BAP) and Chlorpromazine (CPZ) were the positive control chemicals, and potassium chloride (KCl) was the negative control. Each dataset had a total of 10 non-zero dose groups with 3 replicates each and one untreated control group with 12 replicates – we refer the reader to Ramaiahgari et al. [33] for further details. We evaluate the positive control compounds with and without cytotoxic doses, i.e., “all doses” (AD) and “cytotoxic doses removed” (CDR), respectively. Gene expression was measured using the S1500 + TempO-Seq platform [33,37,38]. All datasets’ raw RNA read counts are normalized to a log counts per million scale ( $\log_2(CPM + 1)$ ). Doses were transformed to a  $\log_{10}(dose + 1)$  scale for approximately equally spaced dose groups. The probes that did not meet a minimum read count (MRC), i.e., median  $\log_2(CPM + 1)$  response

equal to 0 and variance less than 0.1, were removed from further analysis.

#### Application of BMD modeling

We fit each probe passing the MRC criteria individually as described above. For the BLX fits, a total of 50,000 MCMC iterations were taken with the first 10,000 disregarded as burn-in, which are the values used in [32]. Trace plots were analyzed and showed quick convergence (usually within 200 iterations) and reasonable mixing. For parametric models, all modeling is performed with BMDEExpress (version 2.3) [6–8].

Estimates of the DR are used to compute a benchmark dose (BMD) for each probe. BMD estimates for both up and down-regulation are computed by solving for BMD in:

$$BMR = \pm \frac{f(BMD) - f(0)}{\sqrt{\frac{1}{\tau}}} \quad (4)$$

Here,  $\tau^{-1}$  is estimated variance, BMR is the benchmark response,  $f(0)$  is the response at the control dose [39]. In this definition, the BMR indicates a standard deviation change in gene expression from the control. For our application, we consider  $BMR = 1$ , which is consistent with previous analyses with these datasets Ramaiahgari et al. [33] and NTP et al. [9].

For monotone curves, we compute the BMD using (4). Most parametric models in BMDEExpress assume expression responses are monotone. In the non-monotone cases, the direction of adversity is determined using a linear trend test [6].

For BLX estimates, it is possible an up- and down-regulated BMDs exist. In these cases, we compute a BMD for both directions of adversity, where the curve's shape determines an adverse response. For DR curves having two extrema, we consider the expression changes in the first "bump" to be an artifact of the fitting procedure. Consequently, we choose the largest BMD found after the maximum as the estimate. For DR curves that are 'J' or umbrella-shaped, there may be two solutions for the BMD. In these cases, the smallest of the estimates is chosen to be the BMD estimate. When no solution to (4) exists in the dosing interval, we set the BMD to be one plus the maximum experimental dose (i.e.  $x_D + 1$ ) to place the BMD above the maximum experimental dose. If up- and down-regulated BMDs exist, we use the smallest of these BMDs as the estimate.

Confidence limits in the BMD are estimated relative to the estimation method. For the BLX spline, calculated using MCMC, we take the 2.5<sup>th</sup> and 97.5<sup>th</sup> quantiles of the BMD distribution to estimate the BMDL and BMDU, respectively. For parametric models, estimated using maximum likelihood, the limits are estimated using the profile likelihood. Point and interval estimates of each BMD are transformed from the log-scale to the original dose scale before calculating the ratio of the upper and lower confidence limits (BMDU/BMDL). The BMDU/BMDL ratio is a measure of reliability in the BMD estimate of each probe. Smaller BMDU/BMDL estimates indicate increased certainty in BMD estimates and more reliable fits. We define probes with reliable expression activity as those having a (i) BMD less than the maximum dose and (ii) 95% confidence limit BMDU/BMDL less than 40, which are consistent with the reliability criteria proposed by NTP et al. [9]. When fitting parametric models, the best model fits must also have a goodness-of-fit  $p$ -value  $> 0.1$ , which indicates no significant difference between the observed expression values and those estimated by the model. Probes meeting this criterion have reliable expression activity and remain in the analysis. This is consistent with NTP et al. [9].

#### Gene Clustering

After model fitting, model data from reliable probes are used to obtain the inputs, described in Section 2.2, to the CSIMM clustering algorithm.

In the TempOSeq platform, multiple probes target one gene designated by a common EntrezID, but their targets may be different gene

isoforms displaying different co-expression patterns. Consequently, we do not average these probes' responses, and we remove probes without a designated EntrezID or multiple EntrezIDs.

For CSIMM, the MCMC sampler is run for 50,000 iterations, with the first 10,000 samples disregarded as burn-in. The clustering strategies are assessed by their ability to separate DR trends into co-expression clusters, where separation is assessed using a heat map. Biological enrichment of clustering results is analyzed using Cluster Enrichment Analysis (CLEAN) [40]. This approach compares the level of functional coherence in clusters based on previously curated biological pathway lists and potency estimates for enriched gene sets.

To compute CLEAN scores, we use the human Hallmark Gene List and the Computational Cancer Gene Neighborhood Gene List (C4:CGN or CGN) from the Molecular Signatures Database (MSigDB) to evaluate the enrichment of gene clusters [14–16,41]. Gene sets containing at least three genes are used for the enrichment analyses consistent with NTP et al. [9]. To compute CLEAN scores, we construct a 2x2 table, as described in Freudenberger et al. [40]. Here, rows indicated membership (i.e. yes/no) to an ALOHA cluster, and the columns represent membership (i.e. yes/no) to an MSigDB gene set. A gene set is enriched for a cluster if the Fischer's exact test  $p$ -value [42], controlled using a Benjamin–Hochberg False Discovery Rate adjustment [43], computed from the 2x2 table is less than 0.01. We call this the Fischer's exact test false discovery rate (FFDR). Gene clusters with the smallest FFDR  $p$ -value for at least one enriched gene are considered to be the most relevant. For biological interpretation, based on the criteria from NTP et al. [9], the most relevant cluster representing it must contain at the minimum 3 of its genes and populate the set by at least 5%. Though we include the NTP et al. [9] criteria for defining active pathways since we include the FFDR  $p$ -value threshold criteria for the clustering approaches, our criteria are fundamentally different from those used in Phillips et al. [7], Yang et al. [8], NTP et al. [9].

Potency estimates for enriched gene sets are obtained using the average BMD estimates from the core genes representing each enriched gene set. We define the core genes to be those in the most relevant cluster for an enriched gene set belonging to that set. For parametric fits, potency estimates are the median BMD estimate of the enriched gene sets. For the clustering results on the normalized expression data and the BLX spline fits we use the posterior mean BMD estimate for each gene. We obtain the lower and upper bounds on the gene set BMD – BMDL and BMDU, respectively – using the 0.025 and the 0.975 quantiles of the best BMD estimates for core genes to obtain a 95% credible interval.

#### BMDEExpress Functional Classification Analysis

After parametric model fitting and post-model filtering of probes, we performed functional classification analyses in BMDEExpress using the human Hallmark Gene List from the Molecular Signatures Database (MSigDB) to compare gene set level BMD estimates with those from each of the clustering approaches [6–8,14–16,41]. It should be noted, prior to the functional classification analysis, BMDEExpress consolidates multiple probes mapping to the same EntrezID. The BMD of these probes are averaged together to obtain one BMD value representing their common EntrezID [6–8].

From the pathway enrichment results from BMDEExpress we obtain a gene set level potency estimate for each of the gene sets. For our analysis we use the median BMDL, median BMD, and median BMDU as estimates of the potency estimates for each of these gene sets. We compare these estimates with those obtained from the most biologically relevant cluster for each of the enriched Hallmark gene sets found with various clustering approaches.



## Results

### Model Fitting

#### Model Predictions

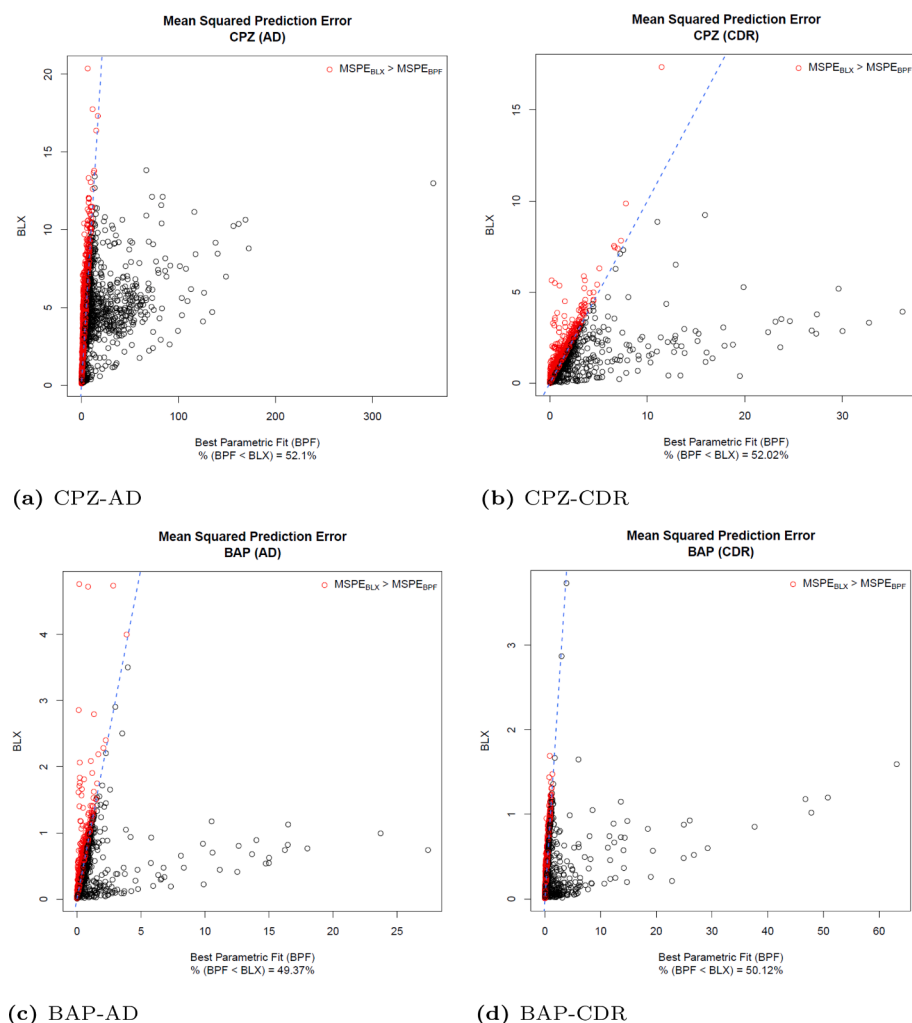
We evaluate each modeling approach's ability to obtain the best estimate of the dose–response curve using a LODOCV prediction. Fig. 2 shows the mean squared prediction error (MSPE) for the parametric and BLX fits for each probe attaining the minimum response criteria in Section 2.3.1 and having less than 75% of their observed responses equal to zero. A smaller MSPE indicates a better prediction. In this figure, the blue dashed line shows situations where the MSPE is the same for a given probe. Black dots, or points to the right of the line, show results where ALOHA outperforms the parametric approach. Red dots, or points to the left of the line, show results where the best parametric model is superior. Here, most of the red points are close to the blue line, which indicates the method fit and predictions are essentially the same for that probe; however, the black points show a greater right skew. ALOHA does a better job describing the dose–responses for these probes, which occurs for approximately 50% of the probes. Further, we observe in general the maximum of the BLX spline MSPEs is an order of magnitude smaller than the maximum of the MSPEs for the best parametric fit, which gives credence to the argument that ALOHA is better representing the true underlying dose–response curve.

Fig. 3 shows why ALOHA is doing better modeling these data and compares the curve fit for a probe from Benzo-[a]-pyrene. The left figure has all doses and while the right one has the highest cytotoxic dose

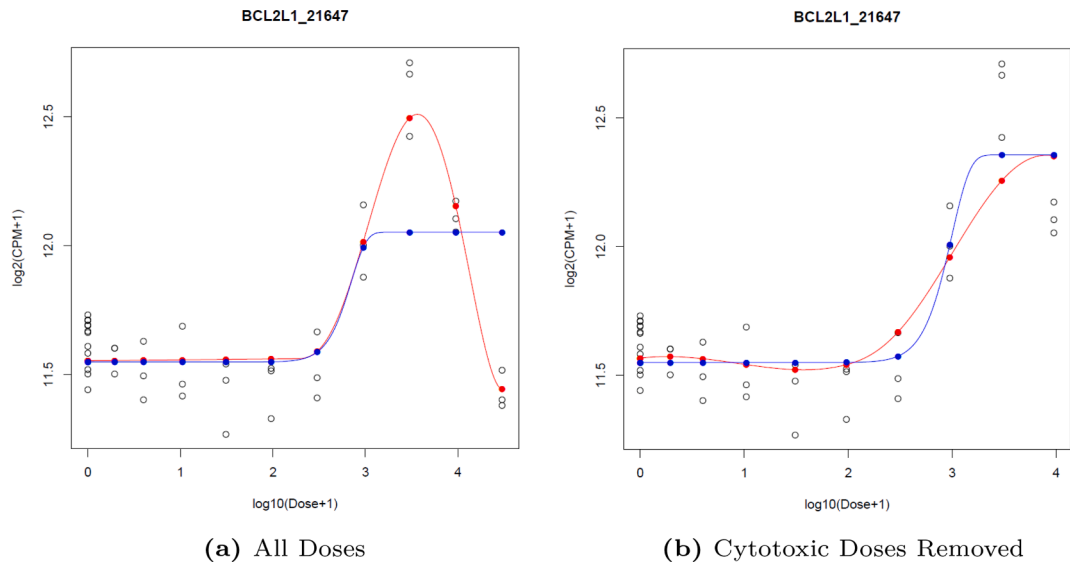
removed. In the left plot, we observe a clear non-monotonic response which the BLX spline fit captures; however, the Exponential 5 model misses this behavior. The right plot shows the fits where the highest dose group is removed since it is cytotoxic. Though this dose's removal is to mitigate issues of non-monotone relationships, there is still a roughly non-monotonic shape. In both of these cases, the model constraints of the best parametric fit, i.e. the Exponential 5, induces a plateau in the high doses, which is unable to capture the peak response and the change in response resulting in a poor fit to the observed data (i.e. fit p-value =  $1e-4$  for both Benzo-[a]-pyrene with and without cytotoxic doses). On the other hand, our approach's flexibility captures the non-monotonic response when all doses are included, and it splits the difference between the response in the two highest dose groups when cytotoxic doses are excluded. There is no longer sufficient evidence to determine if there is a true non-monotonic response.

#### Probe Retention

Out of 2977 probes in the S1500 + TempO-Seq platform, there were 2925, 2703, 2684, 2928, and 2922 probes used in the analysis for Potassium chloride, Benzo-[a]-pyrene with and without cytotoxic doses, and Chlorpromazine with and without cytotoxic doses, respectively, based on MRC filtering described in Section 2.3.1. Table 2 shows the number of modeled probes passing the criteria described in Section 2.3.2 with the two modeling approaches. Other than Benzo-[a]-pyrene with all doses and Chlorpromazine without cytotoxic doses, more probes passed with parametric modeling compared to BLX spline modeling. Many of the same probes were retained between the two modeling



**Fig. 2.** Scatterplots with the mean squared prediction error (MSPE) for the best parametric model fit (BPF) (x-axis) and the BLX spline (y-axis) for each MRC filtered probe modeled with both methods and having less than 75% of observations with a response equal to zero. The blue dashed line indicates where the BLX and BPF MSPE are equivalent. Red points indicate probes with a smaller BPF MSPE than the BLX MSPE. Black points indicate probes with a BLX MSPE smaller than or equal to the BPF MSPE.



**Fig. 3.** Probe passing the reliability criteria with BLX spline model fitting but fails with BMDEExpress (version 2.3) parametric modeling for Benzo-[ $\alpha$ ]-pyrene all doses (left plot) and cytotoxic doses removed (right plot). In both plots the curves display the BLX spline (red) and best parametric model fit from BMDEExpress (blue).

**Table 2**

Number of modeled probes passing model-fit criteria in BMDEExpress (parametric models) and ALOHA (BLX spline model), respectively. Columns 2 & 3 provide the number of probes passing the criteria in Section 2.3.2 for each dataset. Column 4 provides the number of probes in common between the modeling approaches. [AD - all doses; CDR - cytotoxic doses removed].

Compound	BMDEExpress	ALOHA	Shared
Potassium chloride	56	0	0
Benzo[ $\alpha$ ]pyrene (AD)	877	954	464
Benzo[ $\alpha$ ]pyrene (CDR)	973	719	413
Chlorpromazine (AD)	1168	1065	570
Chlorpromazine (CDR)	849	862	362

approaches.

### Gene Clustering Comparison

#### Separation of Dose–Response Curve Shapes

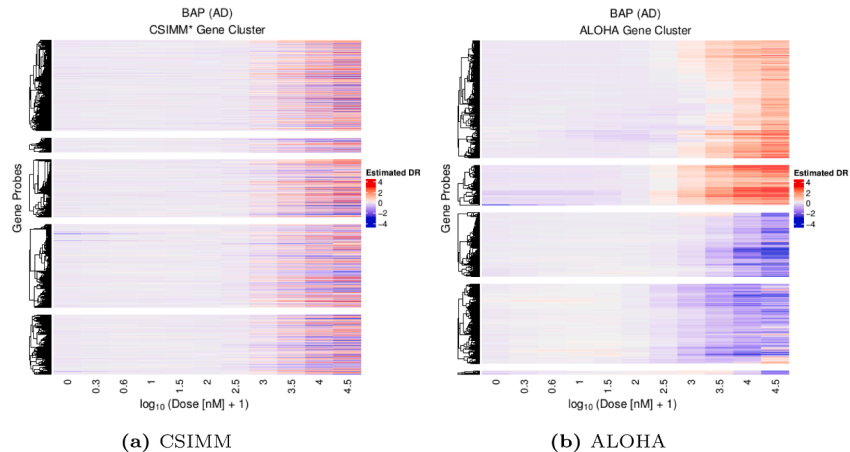
For both Benzo-[ $\alpha$ ]-pyrene and Chlorpromazine, normalized expression values for up and down-regulated genes were better separated with the CSIMM (see S1 Figs19–22). In this case, the BLX spline

parameter estimates showed poor separation, which implies dissimilarity in dose–response curves for genes in the same cluster. For the ALOHA approach, the BLX parameter estimates were separated, which indicates greater similarity in the shape of dose–response curves for genes belonging to the same.

Separation of BLX estimated dose–response data were also evaluated to confirm the coherence of dose–response trends of genes in the same cluster. We chose to cut the posterior pairwise dendrogram such that five representative clusters resulted. The chosen cut-off was arbitrary, but was done for exposition of dose–response trend coherence within clusters. For both Benzo-[ $\alpha$ ]-pyrene and Chlorpromazine, the ALOHA clusters had a high level of coherence in the dose–response patterns compared to CSIMM (see Fig. 4 & S1 Figs23–26). We also calculated several internal validation scores for the five resulting clusters using the Euclidean distance as the measure of dissimilarity [44–47] (see Table 3 & S1 Table 2–5). All validity scores were consistent with greater coherence in dose–response trends of ALOHA gene clusters compared with CSIMM. These results indicate ALOHA is better suited for identifying gene groups with co-regulated dose–response patterns.

#### Functional Coherence in Clusters

CLEAN scores [40] were used to compare the quality of functional



**Fig. 4.** BAP AD - Heatmaps displaying clustering results applied to BLX estimated dose–response data showing visual separation of DR trends. (Left - CSIMM) CSIMM clustering on observed normalized expression data with all replicates in each dose group. (Right - ALOHA) CSIMM clustering on BLX estimated dose–response data.

**Table 3**

Scores measuring the internal coherence of dose–response trends in the five clusters resulting from ALOHA and CSIMM in the Benzo-[a]-pyrene dataset with all doses, see Fig. 4. Larger values for average distance between clusters, Dunn index, and the average silhouette width indicate better separation of genes into similar dose–response groups. Smaller values for the average distance within clusters and entropy indicate more coherence between genes clustered together.

Benzo-[a]-pyrene (AD)	ALOHA	CSIMM
Avg. Distance Between Clusters	3.84	3.23
Avg. Distance Within Clusters	1.68	3.21
Dunn Index	4.4e−3	9.3e−4
Avg. Silhouette Width	0.20	−0.06
Entropy	1.37	1.49

coherence in the resulting co-expression clusters. Fig. 5 compares the empirical cumulative distribution function (ECDF) of CLEAN scores for all clustering approaches on Benzo-[a]-pyrene with all doses. For BMDEExpress clustering, there were fewer genes with a non-zero CLEAN score. Further, the CLEAN scores were generally smaller, e.g., the Benzo-[a]-pyrene with all doses results showed no biological significance with BMDEExpress clustering. Generally, BMDEExpress was the worst of the three approaches at identifying biologically relevant clusters of the three approaches (see S1 Figs27–30). There was little separation between the ECDF curves for CSIMM (black) and ALOHA (red), except in the case of Chlorpromazine with all doses. In that case, there was no clear delineation between any of the curves, which indicates that both approaches obtained similar levels of functional coherence in co-expression clusters (see S1 Figs27–30). Enrichment scores from the most relevant cluster of the enriched gene set showed biological relevance of the clustering approaches was similar (see S1 Figs31–38). Although CSIMM and ALOHA had similar functional coherence, as seen below, CSIMM did not produce clear co-regulated dose–response functions, i.e., the dose responses for genes in a given cluster were not similar (see Fig. 4).

#### Benchmark Dose (BMD) Coherence

Fig. 6 shows the BMD estimate, and the associated 95% credible interval estimate for each enriched Hallmark gene set from the clustering approaches. For enriched gene sets shared between CSIMM and ALOHA, the gene set level BMD estimates were similar. There were a few exceptions to this behavior. For example, in gene set “e2f targets”, the ALOHA cluster BMD interval was smaller, indicating improved BMD coherence. Generally, gene set BMD intervals for BMDEExpress clustering were wider than those from ALOHA.

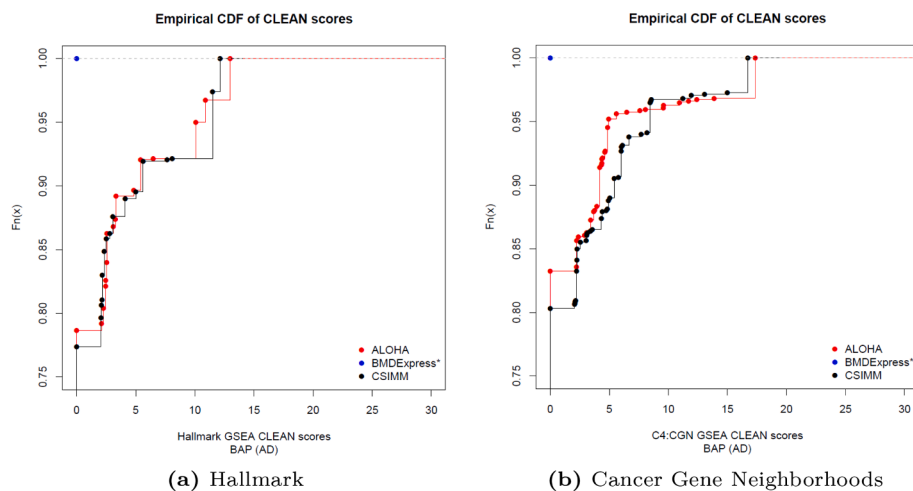
Similar trends were observed for the enriched Cancer Gene Neighborhood gene sets shared between the clustering approaches, except for the Chlorpromazine with all doses which had tighter BMD intervals for many of the gene sets enriched with BMDEExpress clustering (see S1 Figs39–41). Though there were similar BMD estimates between CSIMM and ALOHA, the gene clusters from CSIMM were not representative of a particular dose–response pattern. As a result, BMD estimates were inconsistent about the direction of regulation, and they provided contradictory information about the adverse direction.

Gene set BMD estimates from the clustering approaches were also compared with those estimated after functional classification in BMDEExpress described in Section 2.3.4. Overall, these comparisons showed more genes contributed to the BMD estimate from the current BMDEExpress approach (see Fig. 7). With ALOHA and the other clustering approaches, individual gene BMDs were tighter grouped than those from the current BMDEExpress approach. Confidence intervals for gene set BMDs with the clustering approaches also overlapped or contained those intervals from the current BMDEExpress approach, which indicated clustering resulted in gene groups that produced BMD estimates consistent with BMDEExpress.

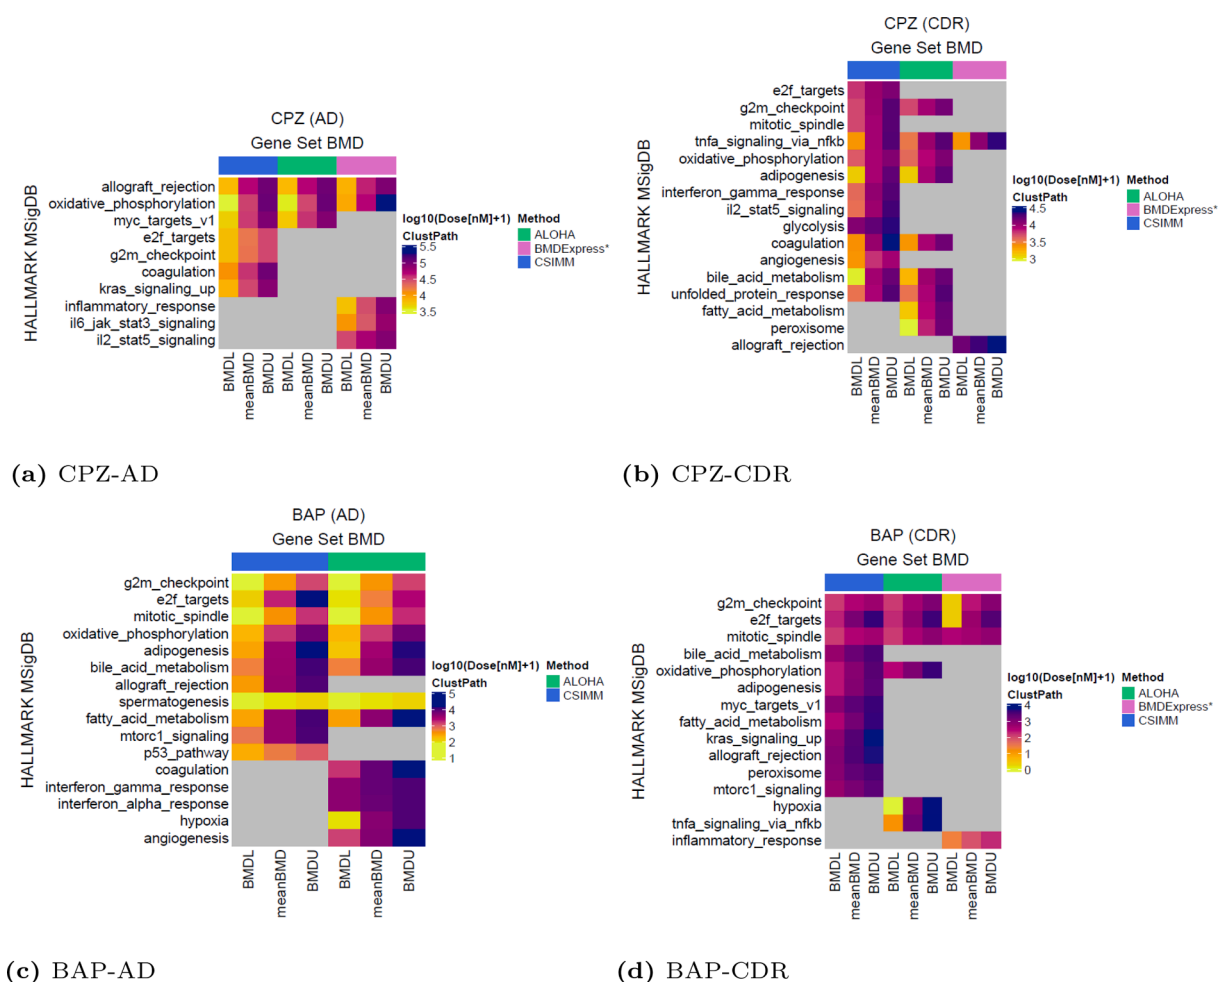
#### Discussion

The use of genomic data in risk assessment relies on developing new methods that account for biological phenomena affecting the underlying data structure. Here, we developed a genomic dose–response (GDR) modeling approach, called Aggregated Local Extrema Splines for High-throughput Dose–Response Analysis (ALOHA), that utilizes model-fit data from Bayesian Local Extremum (BLX) splines and context-specific infinite mixture models (CSIMM) to obtain gene clusters with similar dose–response curves.

For ALOHA, the increased number of probes with a minimum read count (MRC) filtering did not impact the number of probes deemed to have dose–response (DR) information. For example, few probes were removed in the Potassium chloride dataset before fitting, but after DR modeling, there were no remaining probes (see Table 2). This result corresponds to Potassium chloride being a negative control. For probes passing only with parametric modeling, extreme response observations impacted the fit, resulting in an unrealistic fit to the data but an acceptable benchmark dose (BMD) estimate. Here, these observations did not impact the BLX fit, which resulted in greater uncertainty in the BMD estimate or a flat dose–response curve. For probes passing only with BLX, the model flexibility captured non-monotone responses in high doses where cytotoxicity may occur (see Fig. 3). In these cases,



**Fig. 5.** CLEAN score empirical cumulative distribution function (ECDF) plots for Benzo-[a]-pyrene including all doses (BAP-AD), with MSigDB’s Hallmark (right) and Cancer Gene Neighborhood (CGN) (left) gene lists on the normalized expression (CSIMM – black), BLX estimated expression (ALOHA – red), and best parametric fit estimated expression (BMDEExpress – blue) approaches.



**Fig. 6.** Heatmaps comparing the posterior mean BMD and 95% credible intervals (i.e. BMDL and BMDU) estimates (columns) for enriched Hallmark gene sets [14–16] (rows) across the three clustering methods. The BMDs are estimated using core genes in the most relevant cluster, as defined in Section 2.3.3, mapping to the respective gene sets. (a) Chlorpromazine with all experimental dose groups included, (b) Chlorpromazine with cytotoxic dose groups removed, (c) Benzo-[α]-pyrene with all experimental dose groups, (d) Benzo-[α]-pyrene with cytotoxic dose groups removed. Gray blocks within the heatmap indicate which gene sets are not enriched for a clustering approach. Bars on above the columns indicate the various clustering approaches, i.e. CSIMM (blue), ALOHA (green), and BMDEExpress\* clustering (magenta), described in Section 2.2.

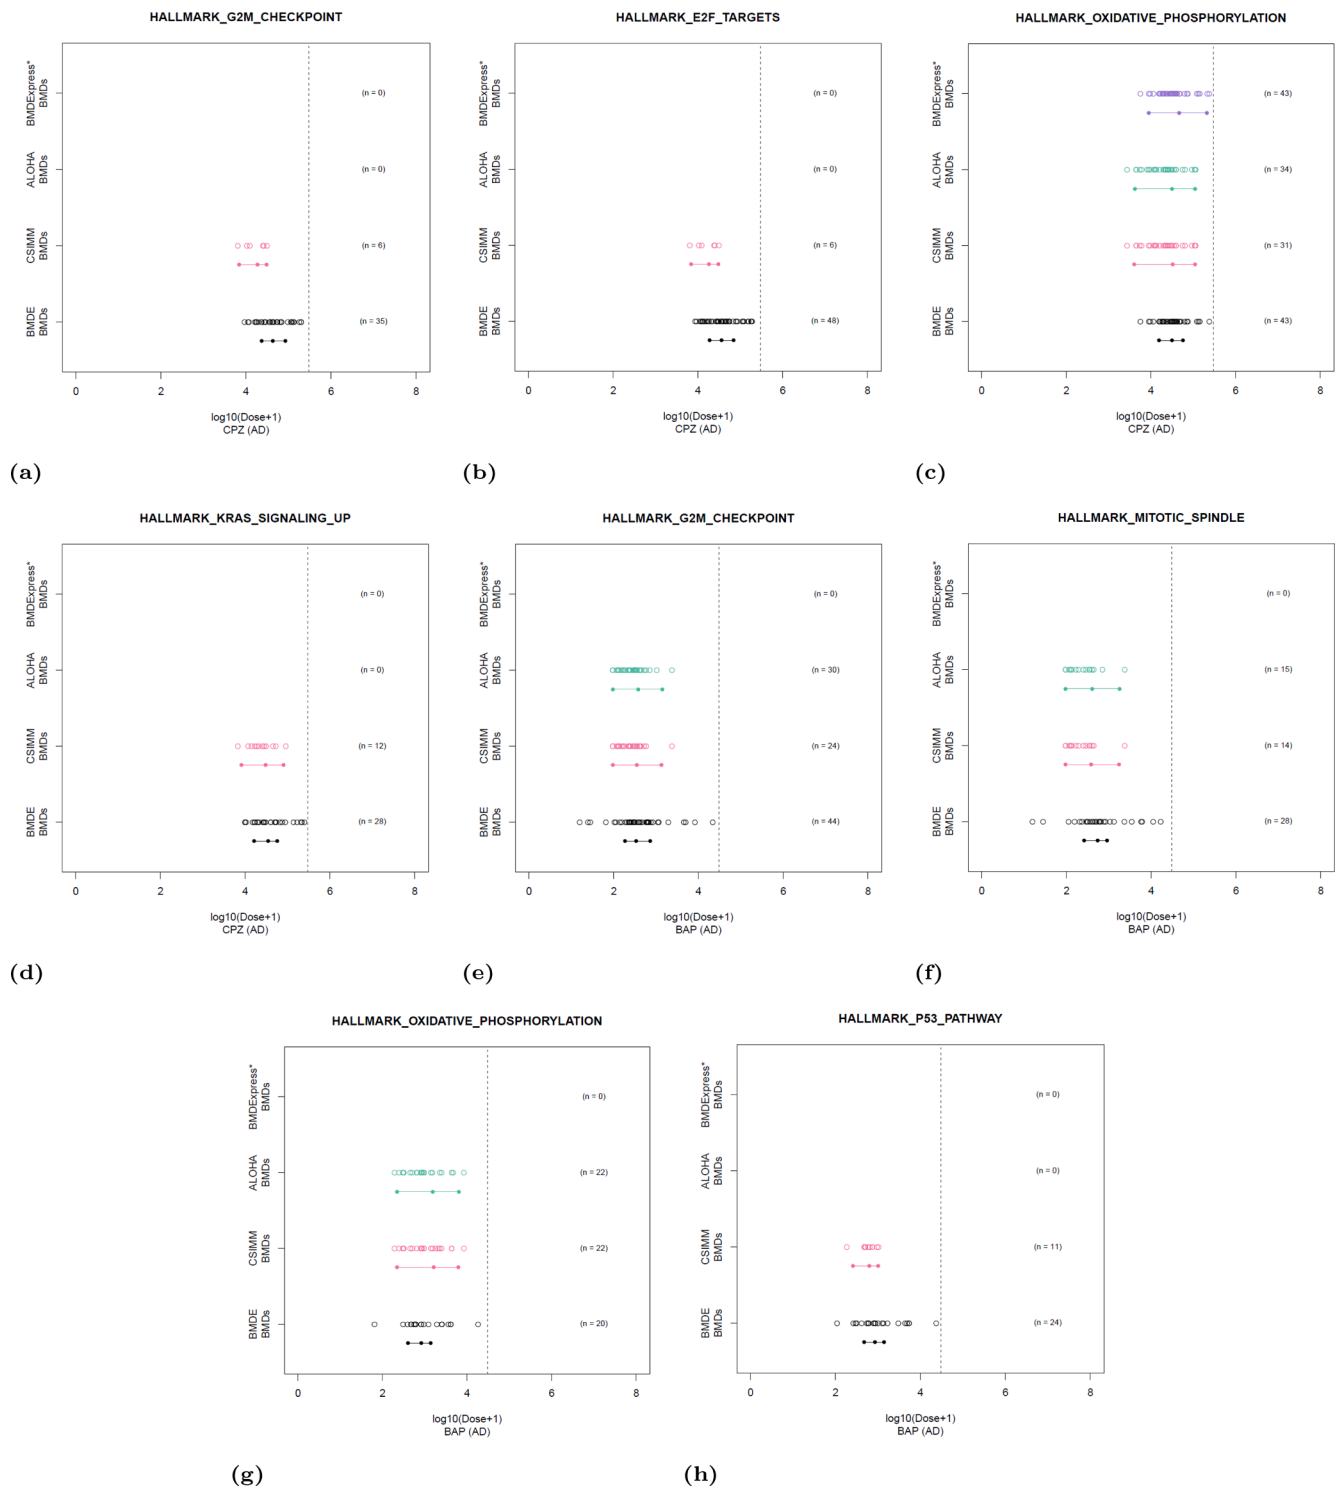
parametric models were unable to fit these observations due to model constraints, resulting in poor fit estimates.

Comparing our method's ability to describe the genomic dose–response data with BMDEExpress (version 2.3), using a LODOCV method, demonstrated little difference between the two approaches in the case of monotone dose–response curves. However, in cases where cytotoxicity induced a non-monotonic response, ALOHA produced smaller mean squared predicted error (MSPE) estimates, thus outperforming the predictive ability of BMDEExpress (see Fig. 2). For many of these genes, the parametric models were unable to appropriately capture the observed response at high doses due to the assumption of monotonicity and model constraints. This assumption led to poor model fits and removal from the analysis. When cytotoxicity is present, a common approach to prevent discarding genes from the analysis is to remove the doses producing the cytotoxic behavior. Though this may resolve some of the genes with non-monotone behavior, it does not guarantee this across all genes. Fig. 3 illustrates this case, where removing the high doses in the Benzo-[α]-pyrene data set did not remove the downturn for “BCL2L1\_21647”. Since the identification of cytotoxic dose groups is an arbitrary decision made by the modeler, any removal approach may result in responsive genes being arbitrarily removed from the analysis. ALOHA does not suffer from this difficulty and retains non-monotone dose–responses for further analyses.

Except for Benzo-[α]-pyrene and Chlorpromazine, with and without cytotoxic doses respectively, more probes passed the post-model fit criteria with BMDEExpress compared to ALOHA. Investigation of individual probe fits showed BMDEExpress inappropriately identified a portion of these as active (i.e., distinct and reliable DR trend). Here, probes fell into one of two categories. First, observed responses were noisy with a few extreme observations (see S1 Fig. 10). Second, the response in high doses was split between active and dead cells (i.e. zero expression) – (see S1 Fig. 11). In the first case, ALOHA estimated a flat curve. In the second case, ALOHA split the difference between active and dead samples' responses in each dose group. The cytotoxic dose range uncertainty did not provide enough evidence to classify the gene as active; either the curve was flat, or there was uncertainty in the BMD estimate near the maximum experimental dose. In both of these cases, constraints in the parametric models were influenced by the most extreme response values from sensitive samples, which produced a curve resulting in a BMD estimate (see examples in S1 Figs. 10 & 11). Whereas ALOHA was robust to sensitive samples, and the overall fit was unaffected without more evidence of an extreme change in response.

ALOHA's ability to filter active genes is shown with potassium chloride, a negative control compound. ALOHA removed all probes from the analysis, unlike BMDEExpress, which retained several genes as active. Some probes in the positive control compound datasets, i.e. Benzo-





**Fig. 7.** Gene set level BMD intervals from BMDExpress (BMDE), using functional classification without clustering, and the clustering approaches on Chlorpromazine and Benzo-[a]-pyrene with all doses (top and bottom, respectively). Points show the best BMD estimate of individual core probes and the number of core probes displayed on the right (see S1 Table 6–13 for details about these probes). (Note: For BMDExpress, probes mapping to the same EntrezID had their BMD estimates averaged and share a single BMD value in the results, but we consider these probes individually in the count to be consistent with core probe counts with clustering.). The vertical dashed line indicates the maximum experimental dose and the line segments below the individual BMDs display the interval (i.e. BMDL to BMDU) and central tendency for the gene set level BMD as described in Section 2.3.3 & Section 2.3.4.

[a]-pyrene and Chlorpromazine, passed the model-fit criteria with BMDExpress but not with ALOHA. Here, the posterior distribution of BMD estimates from ALOHA was either too uncertain or bi-modal with one-mode larger than the maximum dose. In these cases, the 95% credible interval or the posterior mean BMD estimate was unable to pass

the recommended criteria presented in the NTP guidance document [9]. Even without differential gene expression (DEG) filtering and increased data, ALOHA retained active probes (i.e. probes with distinct and reliable DR trends) while removing inactive ones.

Assessing ALOHA's clustering behavior, we first compared the

separation of dose–response patterns in the normalized expression responses and BLX spline parameter estimates – a proxy measure of model fit shapes – between ALOHA and CSIMM [25] (see S1 Fig 19–22). ALOHA showed no clear separation in the regulation trends of normalized expression data. However, there was a clear improvement in the separation of the BLX parameter estimates, which implies DR patterns of individuals' genes are more coherent in ALOHA clusters than CSIMM clusters. To verify these results we compared the separation of the BLX estimated dose–response data, which confirmed the ALOHA clusters contained genes with more coherent DR patterns (see Fig. 4 or S1 Figs. 23 & 26).

Comparison of the functional coherence of the resulting gene clusters demonstrated ALOHA performed similarly to CSIMM. CSIMM did obtain more enriched gene sets than ALOHA. However, the clusters enriched for these gene sets were not necessarily reflective of genes with common DR patterns. Members of a gene set may have different DR curve shapes. Consequently, ALOHA breaks these sets into different clusters, whereas CSIMM clusters these genes regardless if they have conflicting DR patterns. ALOHA also showed clear improvements in the biological relevance of clusters obtained compared to BMDEExpress clustering. These results were likely the result of qualitative differences in probes passed to the clustering process. Alternatively, the limited flexibility of parametric models and inactive probes' retention may have made it difficult for CSIMM to discern clear DR co-expression clusters using BMDEExpress model fit data. In these cases, additional information is needed to distinguish DR patterns better.

Gene set BMD estimates from ALOHA clustering were tighter in most cases than those from BMDEExpress clustering. The only exception was for the Cancer Gene Neighborhood (CGN) gene sets enriched in Chlorpromazine with all doses. This improvement was due to clustering probe fits with more similar DR curves producing consistent BMD estimates. BMD intervals from ALOHA also largely overlapped with intervals estimated using the functional classification results from BMDEExpress, which places all probes passing the model-fit criteria into their respective gene sets. Though BMDEExpress generally produced tighter BMD intervals than ALOHA, the qualitative difference between the two approaches when estimating the confidence intervals was likely the source of this disparity in precision. This result is supported by the observed differences in interval estimates between BMDEExpress clustering and the current BMDEExpress approach. Because both use functional classification, they should be relatively similar (see "OXIDATIVE PHOSPHORYLATION" in see Fig. 7). Individual gene BMDs that contributed to the overall gene set BMD were also more diffuse than ALOHA when using the current BMDEExpress approach. This result shows that even with the addition of clustering, ALOHA was able to obtain BMD estimates consistent with the current approach, BMDEExpress, while identifying gene groups with more coherent BMD estimates.

Genomic dose–response (GDR) analyses currently separate modeling dose-dependent gene expression changes from characterizing coordination among responses. ALOHA integrates these analyses to identify co-expressed genes and characterize altered cellular processes. Our study evaluated three chemical compounds, two eliciting strong transcriptional responses and one with limited to no biological response. Here, ALOHA demonstrated better characterization of GDR relationships, particularly for genes with a non-monotonic response. It did this by retaining more biologically active genes while removing inactive or noisy probes by utilizing shape-constrained splines. The coupling of flexible modeling with co-expression clustering resulted in biologically relevant clusters and potency estimates consistent with the current approaches in BMDEExpress.

ALOHA stream-lines current methods and better accounts for the underlying data structure. Identifying gene clusters with coordinated DR expressions captures cellular processes related to adverse health effects from which gene set level BMDs can be estimated. These estimates may be used to determine a genomic POD comparable to those derived from *in vivo* toxicological endpoints. In summary, ALOHA provides an

alternative step-wise method integrating improved DR models as well as consideration of coordinated expression changes to obtain BMDs that may be used to estimate a biologically relevant POD.

Though we evaluated genomic dose–response data for three compounds, we argue the pipeline applies to similar datasets, even ones with fewer dose groups than those observed here. The primary concern for such a model is that it overfits the data, which the LODOCV analysis showed did not happen. Because fits are reliable dose–response estimates, the remainder of the pipeline, which relies only on the fits' parameters as representations of the dose–response, will produce reliable cluster information on gene co-regulation. The chosen set also represents compounds from varying levels of toxicity, including induction of strong, mild, and minimal transcriptional responses. Consequently, we believe, from a toxicity standpoint, these three chemicals can be seen as generally representative of the range of transcriptional responses possible, which again is an argument for the broader applicability of ALOHA.

While we have demonstrated ALOHA is competitive with the current BMDEExpress, there are several areas for future research. The first is to evaluate the sensitivity and specificity of ALOHA's GDR modeling strategy by including simulated dose–response data. No simulation method to generate this kind of data exists, and developing such a simulation methodology to evaluate ALOHA is an important area of future research. The second would be to optimize the clustering algorithm to obtain disjoint gene clusters that provide a clear delineation for evaluating the accuracy of similar DR patterns in gene clusters and their related biological mechanisms. The last would be to improve the estimation of uncertainty in BMDs from individual and gene set clusters. Such a refinement would allow ALOHA to better quantify the uncertainty between transcriptional and *in vivo* toxicological points-of-departure.

#### CRediT authorship contribution statement

**Sarah E. Davidson:** Conceptualization, Methodology, Software, Validation, Formal analysis, Data curation, Writing - original draft, Writing - review & editing, Visualization, Project administration. **Matthew W. Wheeler:** Conceptualization, Methodology, Software, Validation, Formal analysis, Data curation, Writing - original draft, Writing - review & editing, Visualization, Supervision, Project administration. **Scott S. Auerbach:** Validation, Formal analysis, Data curation, Writing - review & editing. **Siva Sivaganesan:** Validation, Writing - review & editing, Visualization. **Mario Medvedovic:** Conceptualization, Methodology, Validation, Resources, Writing - review & editing, Visualization, Supervision, Project administration, Funding acquisition.

#### Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

#### Acknowledgments

We would like to acknowledge Dr. Davidson's mentor Dr. Eileen D. Kuempel for her guidance and input on investigating the use of toxicogenomics data in occupational health risk assessment, Mr. Nathan M. Drew for his comments on technical issues during initial methods development, Mr. Wen Niu for uploading cleaned expression sets into the iLincs database, the CDC OAMD SCBS team for assistance with improving methods for high-performance computing during pipeline development, and thank Drs. Pierre Bushel and Fred Parham for reviewing the manuscript and their comments.

This work has been supported by the contract from The National Institute for Occupational Safety and Health (75D30118P01991), the grants from National Institutes of Health (U54HL127624 and

P30ES006096), in part by intermural funds at National Institute of Environmental Health Sciences Division of Intramural Research (NIEHS/DIR), and in part by an appointment to the Research Participation Program at the Centers for Disease Control and Prevention administered by the Oak Ridge Institute for Science and Education through an interagency agreement between the U.S. Department of Energy and the Centers for Disease Control and Prevention.

## Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at <https://doi.org/10.1016/j.comtox.2021.100196>.

## References

- [1] R. Farmahin, A. Williams, B. Kuo, N.L. Chepelev, R.S. Thomas, T.S. Barton-Maclaren, I.H. Curran, A. Nong, M.G. Wade, C.L. Yauk, Recommended approaches in the application of toxicogenomics to derive points of departure for chemical risk assessment, *Archives of Toxicology* 91 (5) (2017) 2045–2065, <https://doi.org/10.1007/s00204-016-1886-5>.
- [2] I. Moffat, N.L. Chepelev, S. Labib, J. Bourdon-Lacombe, B. Kuo, J.K. Buick, F. Lemieux, A. Williams, S. Halappanavar, A.I. Malik, et al., Comparison of toxicogenomics and traditional approaches to inform mode of action and points of departure in human health risk assessment of benzo [a] pyrene in drinking water, *Critical Reviews in Toxicology* 45 (1) (2015) 1–43, <https://doi.org/10.1080/10408444.2014.973934>.
- [3] R.S. Thomas, M.A. Philbert, S.S. Auerbach, B.A. Wetmore, M.J. Devito, I. Cote, J.C. Rowlands, M.P. Whelan, S.M. Hays, M.E. Andersen, et al., Incorporating new technologies into toxicity testing and risk assessment: moving from 21st century vision to a data-driven framework, *Toxicological sciences* 136 (1) (2013) 4–18. doi: 10.1093/toxsci/kft178.
- [4] A.F. Webster, N. Chepelev, R. Gagné, B. Kuo, L. Recio, A. Williams, C.L. Yauk, Impact of genomics platform and statistical filtering on transcriptional benchmark doses (bmd) and multiple approaches for selection of chemical point of departure (pod), *PLoS One* 10 (8) (2015), e0136764, <https://doi.org/10.1371/journal.pone.0136764>.
- [5] J.L. Dean, Q.J. Zhao, J.C. Lambert, B.S. Hawkins, R.S. Thomas, S.C. Wesselkamper, Editor's highlight: Application of gene set enrichment analysis for identification of chemically induced, biologically relevant transcriptomic networks and potential utilization in human health risk assessment, *Toxicological Sciences* 157 (1) (2017) 85–99. arXiv:<https://academic.oup.com/toxsci/article-pdf/157/1/85/14173363/kfx021.pdf>, doi:10.1093/toxsci/kfx021. url:<https://doi.org/10.1093/toxsci/kfx021>.
- [6] S.S. Auerbach, *BMDExpress 2.3* (2017). url:<https://github.com/auerbachs/BMDExpress-2/wiki>.
- [7] J.R. Phillips, D.L. Svoboda, A. Tandon, S. Patel, A. Sedykh, D. Mav, B. Kuo, C. L. Yauk, L. Yang, R.S. Thomas, et al., Bmdexpress 2: enhanced transcriptomic dose-response analysis workflow, *Bioinformatics* 35 (10) (2019) 1780–1782, <https://doi.org/10.1093/bioinformatics/bty878>.
- [8] L. Yang, B.C. Allen, R.S. Thomas, Bmdexpress: a software tool for the benchmark dose analyses of genomic data, *BMC Genomics* 8 (1) (2007) 387, <https://doi.org/10.1186/1471-2164-8-387>.
- [9] *NTP RR 5 5* (2473–4756) (2018) 1–44.
- [10] J.-H. Hsieh, A. Sedykh, R. Huang, M. Xia, R.R. Tice, A data analysis pipeline accounting for artifacts in tox21 quantitative high-throughput screening assays, *Journal of Biomolecular Screening* 20 (7) (2015) 887–897, <https://doi.org/10.1177/1087057115581317>.
- [11] U.S. EPA, *Benchmark dose technical guidance document*, Risk Assessment Forum EPA/100/R-12/001, Environmental Protection Agency, U.S., 2012.
- [12] D. Williams, A test for differences between treatment means when several dose levels are compared with a zero dose control, *Biometrics* (1971) 103–117, <https://doi.org/10.2307/2528930>, url:<https://www.jstor.org/stable/2528930>.
- [13] D. Williams, The comparison of several dose levels with a zero dose control, *Biometrics* (1972) 519–531, <https://doi.org/10.2307/2556164>, url:<https://www.jstor.org/stable/2556164>.
- [14] A. Liberzon, A. Subramanian, R. Pinchback, H. Thorvaldsdóttir, P. Tamayo, J.P. Mesirov, Molecular signatures database (msigdb) 3.0, *Bioinformatics* 27 (12) (2011) 1739–1740. arXiv:<https://academic.oup.com/bioinformatics/article-pdf/27/12/1739/716159/btr260.pdf>, doi:10.1093/bioinformatics/btr260. url:<https://doi.org/10.1093/bioinformatics/btr260>.
- [15] A. Liberzon, C. Birger, H. Thorvaldsdóttir, M. Ghandi, J.P. Mesirov, P. Tamayo, The molecular signatures database hallmark gene set collection, *Cell systems* 1 (6) (2015) 417–425, <https://doi.org/10.1016/j.cels.2015.12.004>, url:<https://www.sciencedirect.com/science/article/pii/S2405471215002185>.
- [16] A. Subramanian, P. Tamayo, V.K. Mootha, S. Mukherjee, B.L. Ebert, M.A. Gillette, A. Paulovich, S.L. Pomeroy, T.R. Golub, E.S. Lander, et al., Gene set enrichment analysis: a knowledge-based approach for interpreting genome-wide expression profiles, *Proceedings of the National Academy of Sciences* 102 (43) (2005) 15545–15550, <https://doi.org/10.1073/pnas.0506580102>.
- [17] D.J. Allocchio, I.S. Kohane, A.J. Butte, Quantifying the relationship between co-expression, co-regulation and gene function, *BMC Bioinformatics* 5 (1) (2004) 18, <https://doi.org/10.1186/1471-2105-5-18>.
- [18] D.B. Dahl, *Model-based clustering for expression data via a dirichlet process mixture model*, *Bayesian Inference for Gene Expression and Proteomics* 4 (2006) 201–218.
- [19] C. Fraley, A.E. Raftery, Mclust: Software for model-based cluster analysis, *Journal of Classification* 16 (2) (1999) 297–306. url:<https://EconPapers.repec.org/RePEc:spr:jclass:v:16:y:1999:i:2:p:297-306>.
- [20] M. Medvedovic, S. Sivaganesan, Bayesian infinite mixture model based clustering of gene expression profiles, *Bioinformatics* 18 (9) (2002) 1194–1206, <https://doi.org/10.1093/bioinformatics/18.9.1194>.
- [21] M. Medvedovic, K.Y. Yeung, R.E. Bumgarner, Bayesian mixture model based clustering of replicated microarray data, *Bioinformatics* 20 (8) (2004) 1222–1232, <https://doi.org/10.1093/bioinformatics/bth068>.
- [22] K.Y. Yeung, C. Fraley, A. Murua, A.E. Raftery, W.L. Ruzzo, Model-based clustering and data transformations for gene expression data, *Bioinformatics* 17 (10) (2001) 977–987, <https://doi.org/10.1093/bioinformatics/17.10.977>.
- [23] Z. Bar-Joseph, G. Gerber, D.K. Gifford, T.S. Jaakkola, I. Simon, A new approach to analyzing gene expression time series data, in: *Proceedings of the Sixth Annual International Conference on Computational Biology, RECOMB '02*, Association for Computing Machinery, New York, NY, USA, 2002, p. 39–48. doi:10.1145/565196.565202. url:<https://doi.org/10.1145/565196.565202>.
- [24] P. Kirk, J.E. Griffin, R.S. Savage, Z. Ghahramani, D.L. Wild, Bayesian correlated clustering to integrate multiple datasets, *Bioinformatics* 28 (24) (2012) 3290–3297, <https://doi.org/10.1093/bioinformatics/bts595>.
- [25] X. Liu, S. Sivaganesan, K.Y. Yeung, J. Guo, R.E. Bumgarner, M. Medvedovic, Context-specific infinite mixtures for clustering gene expression profiles across diverse microarray datasets, *Bioinformatics* 22 (14) (2006) 1737–1744, <https://doi.org/10.1093/bioinformatics/btl184>.
- [26] E.F. Lock, D.B. Dunson, Bayesian consensus clustering, *Bioinformatics* 29 (20) (2013) 2610–2616, <https://doi.org/10.1093/bioinformatics/btt425>.
- [27] Y. Luan, H. Li, Clustering of time-course gene expression data using a mixed-effects model with b-splines, *Bioinformatics* 19 (4) (2003) 474–482, <https://doi.org/10.1093/bioinformatics/btg014>.
- [28] Y. Luan, H. Li, Model-based methods for identifying periodically expressed genes based on time course microarray gene expression data, *Bioinformatics* 20 (3) (2004) 332–339, <https://doi.org/10.1093/bioinformatics/btg413>.
- [29] I.C. McDowell, D. Manandhar, C.M. Vockley, A.K. Schmid, T.E. Reddy, B. Engelhardt, Clustering gene expression time series data using an infinite gaussian process mixture model, *PLoS Computational Biology* 14 (1) (2018), e1005896, <https://doi.org/10.1371/journal.pcbi.1005896>.
- [30] J.-H. Park, M. Kyung, Bayesian curve fitting and clustering with dirichlet process mixture models for microarray data, *Journal of the Korean Statistical Society* 48 (2) (2019) 207–220, <https://doi.org/10.1016/j.jkss.2018.11.002>.
- [31] R.S. Savage, Z. Ghahramani, J.E. Griffin, B.J. De la Cruz, D.L. Wild, Discovering transcriptional modules by bayesian data integration, *Bioinformatics* 26 (12) (2010) i158–i167, <https://doi.org/10.1093/bioinformatics/btq210>.
- [32] M.W. Wheeler, D.B. Dunson, A.H. Herring, Bayesian local extremum splines, *Biometrika* 104 (4) (2017) 939–952, <https://doi.org/10.1093/biomet/asx039>.
- [33] S.C. Ramaiahgari, S.S. Auerbach, T.O. Saddler, J.R. Rice, P.E. Dunlap, N.S. Sipes, M.J. DeVito, R.R. Shah, P.R. Bushel, B.A. Merrick, et al., The power of resolution: contextualized understanding of biological responses to liver injury chemicals using high-throughput transcriptomics and benchmark concentration modeling, *Toxicological Sciences* 169 (2) (2019) 553–566, <https://doi.org/10.1093/toxsci/kfz065>.
- [34] H. Akaike, B.N. Petrov, F. Csaki, Second international symposium on information theory (1973).
- [35] H. Akaike, A new look at the statistical model identification, *IEEE Transactions on Automatic Control* 19 (6) (1974) 716–723, <https://doi.org/10.1109/TAC.1974.1100705>.
- [36] Y. Sakamoto, M. Ishiguro, G. Kitagawa, Akaike information criterion statistics, *Dordrecht, The Netherlands: D. Reidel* 81. doi:10.1080/01621459.1988.10478680.
- [37] D. Mav, R.R. Shah, B.E. Howard, S.S. Auerbach, P.R. Bushel, J.B. Collins, D.L. Gerhold, R.S. Judson, A.L. Karmaus, E.A. Maull, et al., A hybrid gene selection approach to create the s1500+ targeted gene sets for use in high-throughput transcriptomics, *PLoS one* 13 (2). doi:10.1371/journal.pone.0191105.
- [38] J.M. Yeakley, P.J. Shepard, D.E. Goyena, H.C. VanSteenhouse, J.D. McComb, B.E. Seligmann, A trichostatin a expression signature identified by tempo-seq targeted whole transcriptome profiling, *PLoS One* 12 (5). doi:10.1371/journal.pone.0178302.
- [39] K.S. Crump, Calculation of benchmark doses from continuous data, *Risk Analysis* 15 (1) (1995) 79–89.
- [40] J.M. Freudenberg, V.K. Joshi, Z. Hu, M. Medvedovic, Clean: Clustering enrichment analysis, *BMC Bioinformatics* 10 (1) (2009) 234, <https://doi.org/10.1186/1471-2105-10-234>.
- [41] A. Liberzon, A description of the molecular signatures database (msigdb) web site, in: *Stem Cell Transcriptional Networks*, Springer, 2014, pp. 153–160.
- [42] S. Lydersen, M.W. Fagerland, P. Laake, Recommended tests for association in 2 × 2 tables, *Statistics in Medicine* 28 (7) (2009) 1159–1175, <https://doi.org/10.1002/sim.3531>.
- [43] Y. Benjamini, Y. Hochberg, Controlling the false discovery rate: a practical and powerful approach to multiple testing, *Journal of the Royal Statistical Society: Series B (Methodological)* 57 (1) (1995) 289–300, <https://doi.org/10.1111/j.2517-6161.1995.tb02031.x>.
- [44] J.C. Dunn, A fuzzy relative of the isodata process and its use in detecting compact well-separated clusters, *Journal of Cybernetics* 3 (3) (1973) 32–57, <https://doi.org/10.1080/01969727308546046>.

- [45] J.C. Dunn, Well-separated clusters and optimal fuzzy partitions, *Journal of Cybernetics* 4 (1) (1974) 95–104, <https://doi.org/10.1080/01969727408546059>.
- [46] M. Meilă, W. Pentney, Clustering by weighted cuts in directed graphs, in: *Proceedings of the 2007 SIAM international conference on data mining*, SIAM, 2007, pp. 135–144, <https://doi.org/10.1137/1.9781611972771.13>.
- [47] P.J. Rousseeuw, Silhouettes: a graphical aid to the interpretation and validation of cluster analysis, *Journal of Computational and Applied Mathematics* 20 (1987) 53–65, [https://doi.org/10.1016/0377-0427\(87\)90125-7](https://doi.org/10.1016/0377-0427(87)90125-7).