



An analytic approach for interpretable predictive models in high-dimensional data in the presence of interactions with exposures

Sahir Rai Bhatnagar^{1,2}  | Yi Yang³ | Budhachandra Khundrakpam⁴ | Alan C. Evans⁴ | Mathieu Blanchette⁵ | Luigi Bouchard⁶ | Celia M.T. Greenwood^{1,2} 

¹Department of Epidemiology, Biostatistics and Occupational Health, McGill University, Montréal, QC, Canada

²Lady Davis Institute, Jewish General Hospital, Montréal, QC, Canada

³Department of Mathematics and Statistics, McGill University, Montréal, QC, Canada

⁴Montreal Neurological Institute, McGill University, Montréal, QC, Canada

⁵Department of Computer Science, McGill University, Montréal, QC, Canada

⁶Department of Biochemistry, Université de Sherbrooke, QC, Canada

Correspondence to:

Celia M.T. Greenwood, Department of Epidemiology, Biostatistics and Occupational Health, McGill University, Montréal, QC, Canada.

Email: celia.greenwood@mcgill.ca

Funding information

Ludmer Centre for Neuroinformatics and Mental Health; Canadian Institutes of Health Research, Grant/Award Number: 363444

ABSTRACT

Predicting a phenotype and understanding which variables improve that prediction are two very challenging and overlapping problems in the analysis of high-dimensional (HD) data such as those arising from genomic and brain imaging studies. It is often believed that the number of truly important predictors is small relative to the total number of variables, making computational approaches to variable selection and dimension reduction extremely important. To reduce dimensionality, commonly used two-step methods first cluster the data in some way, and build models using cluster summaries to predict the phenotype. It is known that important exposure variables can alter correlation patterns between clusters of HD variables, that is, alter network properties of the variables. However, it is not well understood whether such altered clustering is informative in prediction. Here, assuming there is a binary exposure with such network-altering effects, we explore whether the use of exposure-dependent clustering relationships in dimension reduction can improve predictive modeling in a two-step framework. Hence, we propose a modeling framework called ECLUST to test this hypothesis, and evaluate its performance through extensive simulations. With ECLUST, we found improved prediction and variable selection performance compared to methods that do not consider the environment in the clustering step, or to methods that use the original data as features. We further illustrate this modeling framework through the analysis of three data sets from very different fields, each with HD data, a binary exposure, and a phenotype of interest. Our method is available in the *eclust* CRAN package.

KEYWORDS

gene-environment interaction, high-dimensional clustering, prediction models, topological overlap matrix, penalized regression

1 | INTRODUCTION

In this paper, we consider the prediction of an outcome variable y observed on n individuals from p variables, where p is much larger than n . Challenges in this high-dimensional (HD) context include not only building a good predictor which

will perform well in an independent data set, but also being able to interpret the factors that contribute to the predictions. This latter issue can be very challenging in ultra-HD predictor sets. For example, multiple different sets of covariates may provide equivalent measures of goodness of fit (Fan, Han, & Liu, 2014), and therefore how does one decide which are

important? If many variables are highly correlated, interpretation may be improved by acknowledging the existence of an underlying or latent factor generating these patterns. In consequence, many authors have suggested a two-step procedure where the first step is to cluster or group variables in the design matrix in an interpretable way, and then to perform model fitting in the second step using a summary measure of each group of variables.

There are several advantages to these two-step methods. Through the reduction of the dimension of the model, the results are often more stable with smaller prediction variance, and through identification of sets of correlated variables, the resulting clusters can provide an easier route to interpretation. From a practical point of view, two-step approaches are both flexible and easy to implement because efficient algorithms exist for both clustering (e.g., Müllner, 2013) and model fitting (e.g., Friedman, Hastie, & Tibshirani, 2010; Kuhn, 2008; Yang & Zou, 2014), particularly in the case when the outcome variable is continuous.

This two-step idea dates back to 1957 when Kendall first proposed using principal components in regression (Kendall, 1957). Hierarchical clustering based on the correlation of the design matrix has also been used to create groups of genes in microarray studies. For example, at each level of a hierarchy, cluster averages have been used as new sets of potential predictors in both forward-backward selection (Hastie, Tibshirani, Botstein, & Brown, 2001) or the lasso (Park, Hastie, & Tibshirani, 2007). Bühlmann et al. proposed a bottom-up agglomerative clustering algorithm based on canonical correlations and used the group lasso on the derived clusters (Bühlmann, Rütimann, van de Geer, & Zhang, 2013). A more recent proposal performs sparse regression on cluster prototypes (Reid & Tibshirani, 2016), that is, extracting the most representative gene in a cluster instead of averaging them.

These two-step approaches usually group variables based on a matrix of correlations or some transformation of the correlations. However, when there are external factors, such as exposures, that can alter correlation patterns, a dimension reduction step that ignores this information may be suboptimal. Many of the HD genomic data sets currently being generated capture a possibly dynamic view of how a tissue is functioning, and demonstrate differential patterns of coregulation or correlation under different conditions. We illustrate this critical point with an example of a microarray gene expression data set available in the *COPDSexualDimorphism.data* package (Sathirapongsasuti, 2013) from Bioconductor. This study measured gene expression in chronic obstructive pulmonary disease (COPD) patients and controls in addition to their age, gender, and smoking status. To see if there was any effect of smoking status on gene expression, we plotted the expression profiles separately for current and never smokers. To balance the covariate profiles, we matched subjects from

each group on age, gender and COPD case status, resulting in a sample size of 7 in each group. Heatmaps in Figure 1 show gene expression levels and the corresponding gene-gene correlation matrices as a function of dichotomized smoking status for 2,900 genes with large variability. Evidently, there are substantial differences in correlation patterns between the smoking groups (Figures 1a and 1b). However, it is difficult to discern any patterns or major differences between the groups when examining the gene expression levels directly (Figures 1c and 1d). This example highlights two key points: (1) environmental exposures can have a widespread effect on regulatory networks and (2) this effect may be more easily discerned by looking at a measure for gene similarity, relative to analyzing raw expression data.

Many other examples of altered coregulation and phenotype associations can be found. For instance, in a pediatric brain development study, very different correlation patterns of cortical thickness within brain regions were observed across age groups, consistent with a process of fine-tuning an immature brain system into a mature one (Khundrakpam et al., 2013). A comparison of gene expression levels in bone marrow from 327 children with acute leukemia found several differentially *coexpressed* genes in Philadelphia-positive leukemias compared to the cytogenetically normal group (Kostka & Spang, 2004). To give the third example, an analysis of RNA-sequencing data from The Cancer Genome Atlas (TCGA) revealed very different correlation patterns among sets of genes in tumors grouped according to their missense or null mutations in the TP53 tumor suppressor gene (Oros Klein et al., 2016).

Therefore, in this paper, we pose the question whether clustering or dimension reduction that incorporates known covariate or exposure information can improve prediction models in HD genomic data settings. Substantial evidence of dysregulation of genomic coregulation has been observed in a variety of contexts, however we are not aware of any work that carefully examines how this might impact the performance of prediction models. We propose a conceptual analytic strategy called ECLUST, for prediction of a continuous or binary outcome in HD contexts while exploiting exposure-sensitive data clusters. We restrict our attention to two-step algorithms in order to implement a covariate-driven clustering.

Specifically, we hypothesize that within two-step methods, variable grouping that considers exposure information can lead to improved predictive accuracy and interpretability. We use simulations to compare our proposed method to comparable approaches that combine data reduction with predictive modeling. We are focusing our attention primarily on the performance of alternative dimension reduction strategies within the first step of a two-step method. Therefore, performance of each strategy is compared for several appropriate Step 2 predictive models. We then illustrate these concepts more concretely by analyzing three data sets. Our method and the

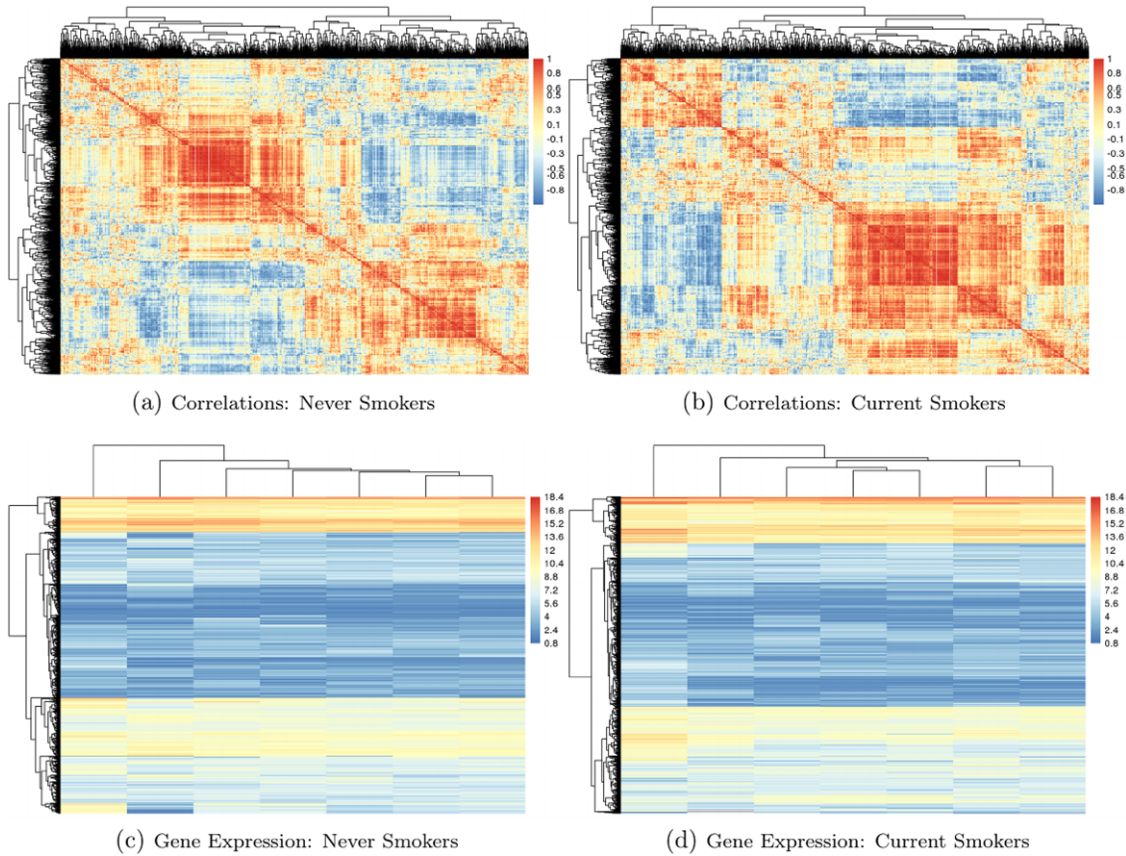


FIGURE 1 Heatmaps of correlations between genes (top) and gene expression data (bottom: rows are genes and columns are subjects) stratified by smoking status from a microarray study of COPD (Sathirapongsasuti, 2013). The 20% most variable genes are displayed (2,900 genes)

Note: There are seven subjects in each group, matched on COPD case status, gender, and age. Data are available at Bioconductor in the COPDSexualDimorphism.data package.

functions used to conduct the simulation studies have been implemented in the R package *eclust* (Bhatnagar, 2017), available on CRAN. Extensive documentation of the package is available at <http://sahirbhatnagar.com/eclust/>.

2 | METHODS

Assume that there is a single binary environmental factor E of importance, and an $n \times p$ HD data set \mathbf{X} (n observations, p features) of relevance. This could be genome-wide epigenetic data, gene expression data, or brain imaging data, for example. Assume that there is a continuous or binary phenotype of interest Y and that the environment has a widespread effect on the HD data, that is, affects many elements of the HD data. The primary goal is to improve the prediction of Y by identifying interactions between E and \mathbf{X} through a carefully constructed data reduction strategy that exploits E -dependent correlation patterns. The secondary goal is to improve identification of the elements of \mathbf{X} that are involved; we denote this subset by S_0 . We hypothesize that a systems-based perspective will be informative when exploring the factors that are

associated with a phenotype of interest, and in particular we hypothesize that incorporation of environmental factors into predictive models in a way that retains a HD perspective will improve results and interpretation.

2.1 | Potential impacts of covariate-dependent coregulation

Motivated by real-world examples of differential coexpression, we first demonstrate that environment-dependent correlations in \mathbf{X} can induce an interaction model. Without loss of generality, let $p = 2$ and the relationship between X_1 and X_2 depend on the environment such that

$$X_{i2} = \psi X_{i1} E_i + \varepsilon_i, \quad (1)$$

where ε_i is an error term and ψ is a slope parameter, that is,

$$X_{i2} = \begin{cases} \psi X_{i1} + \varepsilon_i & \text{when } E_i = 1 \\ \varepsilon_i & \text{when } E_i = 0. \end{cases}$$

Consider the 3-predictor regression model

$$Y_i = \beta_0 + \beta_1 X_{i1} + \beta_2 X_{i2} + \beta_3 E_i + \varepsilon_i^*, \quad (2)$$

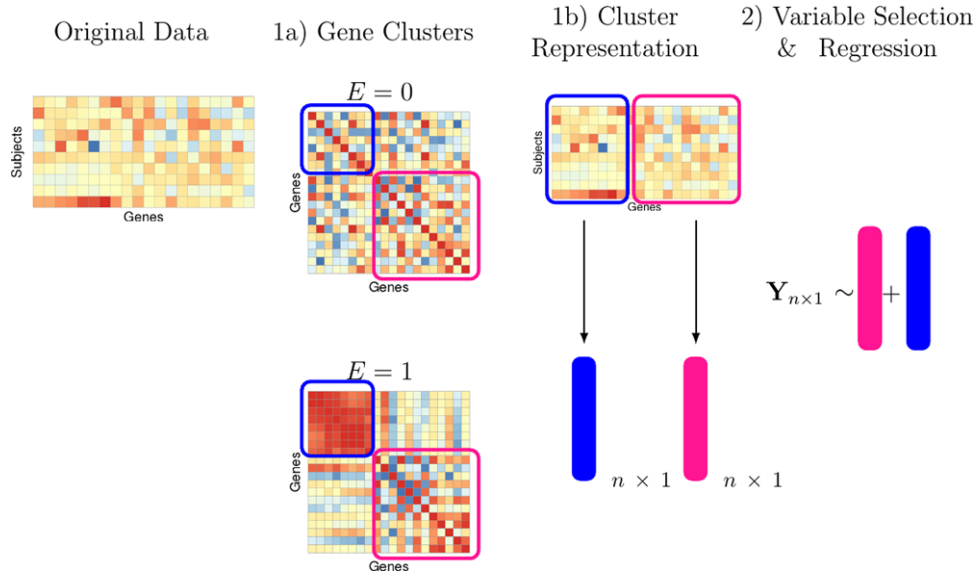


FIGURE 2 An overview of our proposed method

Note: (1a) A measure of similarity is calculated separately for both groups and clustering is performed on a linear combination of these two matrices. (1b) We reduce the dimension of each cluster by taking a summary measure. (2) Variable selection and regression is performed on the cluster representatives, E , and their interaction with E

where ϵ_i^* is another error term which is independent of ϵ_i . At first glance, (2) does not contain any interaction terms. However, substituting (1) for X_{i2} in (2) we get

$$Y_i = \beta_0 + \beta_1 X_{i1} + \beta_2 \psi(X_{i1} E_i) + \beta_3 E_i + \epsilon_i \beta_2 + \epsilon_i^*. \quad (3)$$

The third term in (3) resembles an interaction model, with $\beta_2 \psi$ being the interaction parameter. We present the second illustration showing how nonlinearity can induce interactions. Suppose

$$Y_i = \beta_0 + \beta_1 X_{i1} + \beta_2 X_{i2} + \beta_3 E_i + \beta_3 \max_{j \in \{1,2\}} (X_{ij} - \bar{X}_i)^2 + \epsilon_i^*. \quad (4)$$

Substituting (1) for X_{i2} in (4), we obtain a nonlinear interaction term. Equation (4) provided partial motivation for the model used in our third simulation scenario. Some motivation for this model and a graphical representation are presented in the Section 3.

2.2 | Proposed framework and algorithm

We restrict attention to methods containing two phases as illustrated in Figure 2(1a), a clustering stage where variables are clustered based on some measure of similarity, Figure 2(1b), a dimension reduction stage where a summary measure is created for each of the clusters, and Figure 2(2), a simultaneous variable selection and regression stage on the summarized cluster measures. Although this framework appears very similar to any two-step approach, our hypothesis is that allowing the clustering in Step 1a to depend on

the environment variable can lead to improvements in prediction after Step 2. Hence, methods in Step 1a are adapted to this end, as described in the following sections. Our focus in this manuscript is on the clustering and cluster representation steps. Therefore, we compare several well-known methods for variable selection and regression that are best adapted to our simulation designs and data sets.

2.2.1 | Step 1a: Clustering using coexpression networks that are influenced by the environment

In agglomerative clustering, a measure of similarity between sets of observations is required in order to decide which clusters should be combined. Common choices include Euclidean, maximum, and absolute distance. A more natural choice in genomic or brain imaging data is to use the Pearson correlation (or its absolute value) because the derived clusters are biologically interpretable. Indeed, genes that cluster together are correlated and thus likely to be involved in the same cellular process. Similarly, cortical thickness measures of the brain tend to be correlated within predefined regions such as the left and right hemisphere, or frontal and temporal regions (Sato et al., 2013). However, the information on the connection between two variables, as measured by the Pearson correlation, for example, may be noisy or incomplete. Thus it is of interest to consider alternative measures of pairwise interconnectedness. Gene coexpression networks are being used to explore the system-level function of genes, where nodes represent genes and are connected if they are significantly coexpressed (Zhang & Horvath, 2005), and here we

use their overlap measure (Ravasz, Somera, Mongru, Oltvai, & Barabási, 2002) to capture connectedness between two X variables within each environmental condition. As was discussed earlier, genes can exhibit very different patterns of correlation in one environment versus the other (e.g. Figure 1). Furthermore, measures of similarity that go beyond pairwise correlations and consider the shared connectedness between nodes can be useful in elucidating networks that are biologically meaningful. Therefore, we propose to first look at the topological overlap matrix (TOM) separately for exposed ($E = 1$) and unexposed ($E = 0$) individuals (see supplementary Section A for details on the TOM). We then seek to identify nodes that are very different between environments. We determine differential coexpression using the absolute difference $TOM(X_{\text{diff}}) = |TOM_{E=1} - TOM_{E=0}|$ (Oros Klein et al., 2016). We then use hierarchical clustering with average linkage on the derived difference matrix to identify these differentially coexpressed variables. Clusters are automatically chosen using the dynamicTreeCut (Langfelder, Zhang, & Horvath, 2008) algorithm. Of course, there could be other clusters which are not sensitive to the environment. For this reason, we also create a set of clusters based on the TOM for all subjects denoted $TOM(X_{\text{all}})$. This will lead to each covariate appearing in two clusters. In the sequel, we denote the clusters derived from $TOM(X_{\text{all}})$ as the set $C_{\text{all}} = \{C_1, \dots, C_k\}$, and those derived from $TOM(X_{\text{diff}})$ as the set $C_{\text{diff}} = \{C_{k+1}, \dots, C_\ell\}$, where $k < \ell < p$.

2.2.2 | Step 1b: Dimension reduction via cluster representative

Once the clusters have been identified in phase 1, we proceed to reduce the dimensionality of the overall problem by creating a summary measure for each cluster. A low-dimensional structure, that is, grouping when captured in a regression model, improves predictive performance and facilitates a model's interpretability. We propose to summarize a cluster by a single representative number. Specifically, we chose the average values across all measures (Bühlmann et al., 2013; Park et al., 2007), and the first principal component (Langfelder & Horvath, 2007). These representative measures are indexed by their cluster, that is, the variables to be used in our predictive models are $\tilde{\mathbf{X}}_{\text{all}} = \{\tilde{X}_{C_1}, \dots, \tilde{X}_{C_k}\}$ for clusters that do not consider E , as well as $\tilde{\mathbf{X}}_{\text{diff}} = \{\tilde{X}_{C_{k+1}}, \dots, \tilde{X}_{C_\ell}\}$ for E -derived clusters. The tilde notation on the X is to emphasize that these variables are different from the separate variables in the original data.

2.2.3 | Step 2: Variable selection and regression

Because the clustering in phase 1 is unsupervised, it is possible that the derived latent representations from phase 2 will

not be associated with the response. We therefore use penalized methods for supervised variable selection, including the lasso (Tibshirani, 1996) and elasticnet (Zou & Hastie, 2005) for linear models, and multivariate adaptive regression splines (MARS; Friedman, 1991) for nonlinear models. We argue that the selected nonzero predictors in this model will represent clusters of genes that interact with the environment and are associated with the phenotype. Such an additive model might be insufficient for predicting the outcome. In this case we may directly include the environment variable, the summary measures, and their interaction. In the light of our goals to improve prediction and interpretability, we consider the following model:

$$g(\mu) = \beta_0 + \sum_{j=1}^{\ell} \beta_j \tilde{X}_{C_j} + \beta_E E + \sum_{j=1}^{\ell} \alpha_j (\tilde{X}_{C_j} E), \quad (5)$$

where $g(\cdot)$ is a known link function, $\mu = E[Y|\mathbf{X}, E, \beta, \alpha]$ and \tilde{X}_{C_j} are linear combinations of \mathbf{X} (from Step 1b). The primary comparison is models with $\tilde{\mathbf{X}}_{\text{all}}$ only versus models with $\tilde{\mathbf{X}}_{\text{all}}$ and $\tilde{\mathbf{X}}_{\text{diff}}$. Given the context of either the simulation or the data set, we use either linear models or nonlinear models. Our general approach, ECLUST, can therefore be summarized by the algorithm in Table 1.

3 | SIMULATION STUDIES

We have evaluated the performance of our ECLUST method in a variety of simulated scenarios. For each simulation scenario, we compared ECLUST to the following analytic approaches: (1) regression and variable selection is performed on the model which consists of the original variables, E and their interaction with E (SEPARATE), and (2) clustering is performed without considering the environmental exposure followed by regression and variable selection on the cluster representations, E , and their interaction with E (CLUST). A detailed description of the methods being compared is summarized in Table 2. We have designed six simulation scenarios that illustrate different kinds of relationships between the variables and the response. For all scenarios, we have created HD data sets with p predictors ($p = 5,000$), and sample sizes of $n = 200$. We also assume that we have two data sets for each simulation—a training data set where the parameters are estimated, and a testing data set where prediction performance is evaluated, each of size $n_{\text{train}} = n_{\text{test}} = 200$. The number of subjects who were exposed ($n_{E=1} = 100$) and unexposed ($n_{E=0} = 100$) and the number of truly associated parameters ($|S_0| = 500$) remain fixed across the 6 simulation scenarios. Let

$$Y = Y^* + k \cdot \epsilon, \quad (6)$$

TABLE 1 Details of ECLUST algorithm

Step	Description, Software ^a and Reference
1a)	<p>i) Calculate <i>TOM</i> separately for observations with $E = 0$ and $E = 1$ using <code>WGCNA::TOMsimilarityFromExpr</code> (Langfelder & Horvath, 2008)</p> <p>ii) Euclidean distance matrix of $TOM_{E=1} - TOM_{E=0}$ using <code>stats::dist</code></p> <p>iii) Run the <code>dynamicTreeCut</code> algorithm (Langfelder et al., 2008; Langfelder, P., Zhang, B., & with contributions from Steve Horvath, 2016) on the distance matrix to determine the number of clusters and cluster membership using <code>dynamicTreeCut::cutreeDynamic</code> with <code>minClusterSize = 50</code></p>
1b)	<p>i) 1st PC or average for each cluster using <code>stat::prcomp</code> or <code>base::mean</code></p> <p>ii) Penalized regression model: create a design matrix of the derived cluster representatives and their interactions with E using <code>stats::model.matrix</code></p> <p>iii) MARS model: create a design matrix of the derived cluster representatives and E</p>
2)	<p>i) For linear models, run penalized regression on design matrix from Step 1b using <code>glmnet::cv.glmnet</code> (Friedman et al., 2010). Elasticnet mixing parameter $\alpha=1$ corresponds to the lasso and $\alpha=0.5$ corresponds to the value we used in our simulations for elasticnet. The tuning parameter λ is selected by minimizing 10 fold cross-validated mean squared error (MSE).</p> <p>ii) For nonlinear effects, run MARS on the design matrix from Step 1b using <code>earth::earth</code> (Milborrow. Derived from <code>mda:mars</code> by T. Hastie and R. Tibshirani., 2011) with pruning method <code>pmethod = "backward"</code> and maximum number of model terms <code>nk = 1000</code>. The degree=1,2 is chosen using 10 fold cross validation (CV), and within each fold the number of terms in the model is the one that minimizes the generalized cross validated (GCV) error.</p>

^aAll functions are implemented in R (R Core Team, 2016). The naming convention is as follows: `package_name::package_function`. Default settings used for all functions unless indicated otherwise.

TABLE 2 Summary of methods used in simulation study

General approach	Summary Measure of feature clusters	Description ^{a,b}
SEPARATE	NA	Regression of the original predictors $\{X_1, \dots, X_p\}$ on the response, i.e., no transformation of the predictors is being done here
CLUST	1st PC, average	Create clusters of predictors without using the environment variable $\{C_1, \dots, C_k\}$. Use the summary measure of each cluster as inputs of the regression model.
ECLUST	1st PC, average	Create clusters of predictors using the environment variable $\{C_{k+1}, \dots, C_\ell\}$ where $k < \ell < p$, as well as clusters without the environment variable $\{C_1, \dots, C_k\}$. Use summary measures of $\{C_1, \dots, C_\ell\}$ as inputs of the regression model.

^aSimulations 1 and 2 used lasso and elasticnet for the linear models, and Simulation 3 used MARS for estimating nonlinear effects.

^bSimulations 4–6 convert the continuous response generated in simulations 1–3, respectively, into a binary response.

^cPC: principal component.

where Y^* is the linear predictor, the error term ϵ is generated from a standard normal distribution, and k is chosen such that the signal-to-noise ratio $SNR = (Var(Y^*)/Var(\epsilon))$ is 0.2, 1 and 2 (e.g. the variance of the response variable Y due to ϵ is $1/SNR$ of the variance of Y due to Y^*).

3.1 | The design matrix

We generated covariate data in blocks using the `simulateDatExpr` function from the `WGCNA` package in R (version 1.51). This generates data from a latent vector: first a seed vector is simulated, then covariates are generated with varying degree of correlation with the seed vector in a given block. We simulated five clusters (blocks), each of size 750 variables, and labeled them by color (turquoise, blue, red, green, and yellow), while the remaining 1,250 variables were simulated as independent standard normal vectors (gray) (Figure 3). For the unexposed observations ($E = 0$), only the predictors in the yellow block were simulated with correlation, while all other covariates were independent within and between blocks. The TOM values are very small for the yellow cluster because it is not correlated with any of its neighbors. For the exposed observations ($E = 1$), all five blocks contained predictors that are correlated. The blue and turquoise blocks are set to have an average correlation of 0.6. The average correlation was varied for both green and red clusters $\rho = \{0.2, 0.9\}$ and the active set S_0 , that are directly associated with y , was distributed evenly between these two blocks. Heatmaps of the TOM for this environment-dependent correlation structure are shown in Figure 3 with annotations for the true clusters and active variables. This design matrix shows widespread changes in gene networks in the exposed environment, and this subsequently affects the phenotype through the two associated clusters. There are also pathways that respond to changes in the environment but are not associated with the response (blue and turquoise), while others that are neither active in the disease nor affected by the environment (yellow).

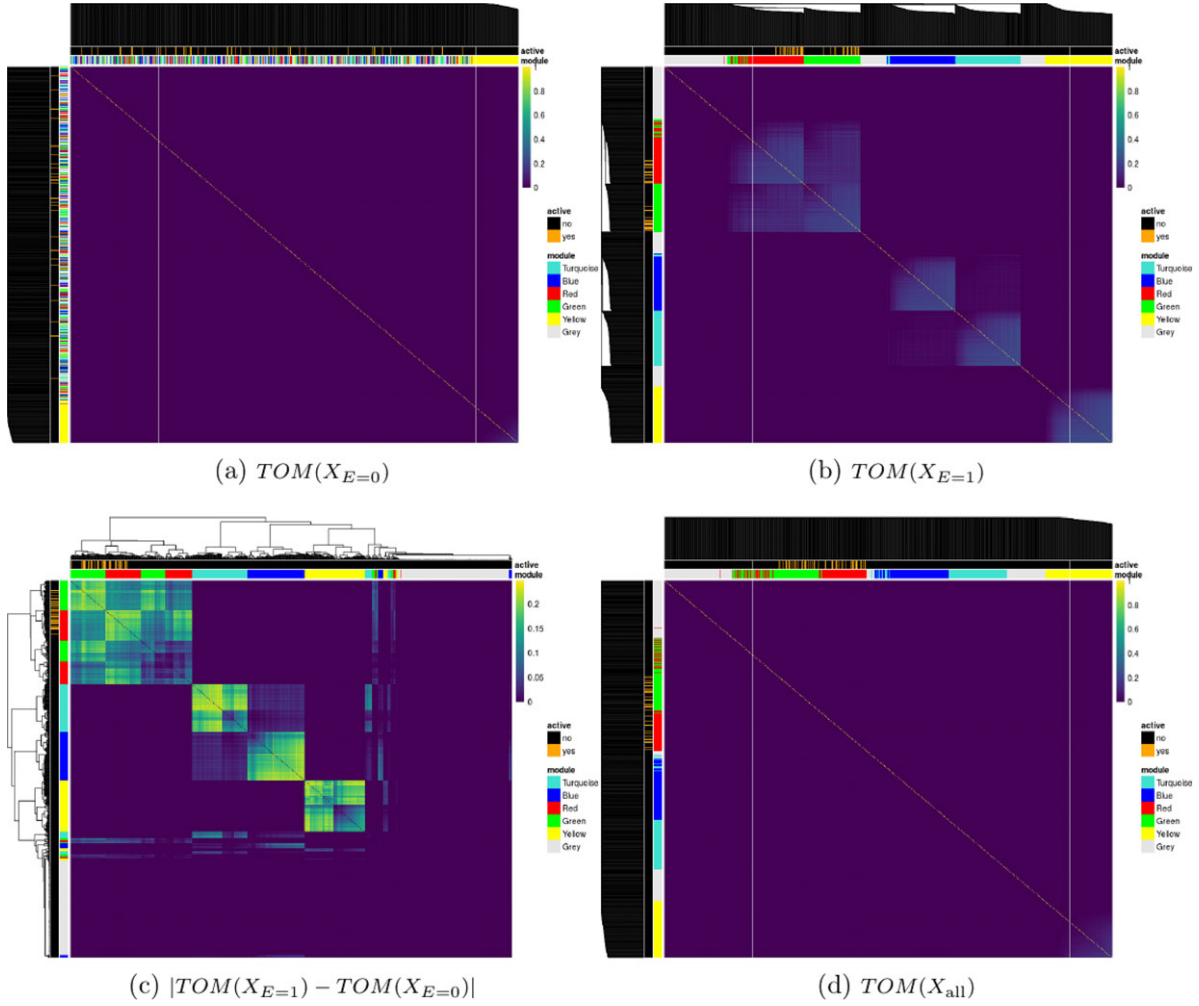


FIGURE 3 Topological overlap matrices (TOM) of simulated predictors based on subjects with (a) $E = 0$, (b) $E = 1$, (c) their absolute difference, and (d) all subjects

Note: Dendrograms are from hierarchical clustering (average linkage) of one minus the TOM for (a), (b), and (d) and the Euclidean distance for (c). Some variables in the red and green clusters are associated with the outcome variable. The *module* annotation represents the true cluster membership for each predictor, and the *active* annotation represents the truly associated predictors with the response.

3.2 | The response

The first three simulation scenarios differ in how the linear predictor Y^* in (6) is defined, and also in the choice of regression model used to fit the data. In simulations 1 and 2, we use lasso (Tibshirani, 1996) and elasticnet (Zou & Hastie, 2005) to fit linear models; then we use MARS (Friedman, 1991) in Simulation 3 to estimate nonlinear effects. In simulations 4–6, we use the GLM version of these models, respectively, because the responses are binary.

3.2.1 | Simulation 1

Simulation 1 was designed to evaluate performance when there are no explicit interactions between X and E (see Equation (3)). We generated the linear predictor from

$$Y^* = \sum_{\substack{j \in \{1, \dots, 250\} \\ j \in \text{red, green block}}} \beta_j X_j + \beta_E E, \quad (7)$$

where $\beta_j \sim \text{Unif}[0.9, 1.1]$ and $\beta_E = 2$. That is, only the first 250 predictors of both red and green blocks are active. In this setting, only the main effects model is being fit to the simulated data.

3.2.2 | Simulation 2

In the second scenario, we explicitly simulated interactions. All nonzero main effects also had a corresponding nonzero interaction effect with E . We generated the linear predictor from

$$Y^* = \sum_{\substack{j \in \{1, \dots, 125\} \\ j \in \text{red, green block}}} (\beta_j X_j + \alpha_j X_j E) + \beta_E E, \quad (8)$$

where $\beta_j \sim \text{Unif}[0.9, 1.1]$, $\alpha_j \sim \text{Unif}[0.4, 0.6]$ or $\alpha_j \sim \text{Unif}[1.9, 2.1]$, and $\beta_E = 2$. In this setting, both main effects and their interactions with E are being fit to the simulated data.




FIGURE 4 Visualization of the relationship between the response, the first principal component of the main effects and $f(Q_i)$ in (9) for $E = 0$ (left) and $E = 1$ (right) in simulation scenario 3

Note: This graphic also depicts the intuition behind model (4).

3.2.3 | Simulation 3

In the third simulation, we investigated the performance of the ECLUST approach in the presence of nonlinear effects of the predictors on the phenotype:

$$Y_i^* = \sum_{\substack{j \in \{1, \dots, 250\} \\ j \in \text{red, green block}}} \beta_j X_{ij} + \beta_E E_i + \alpha_Q E_i \cdot f(Q_i), \quad (9)$$

where

$$Q_i = - \max_{\substack{j \in \{1, \dots, 250\} \\ j \in \text{red, green block}}} (X_{ij} - \bar{X}_i)^2 \quad (10)$$

$$f(u_i) = \frac{u_i - \min_{i \in \{1, \dots, n\}} u_i}{-\min_{i \in \{1, \dots, n\}} u_i} \quad (11)$$

$$\bar{X}_i = \frac{1}{500} \sum_{\substack{j \in \{1, \dots, 250\} \\ j \in \text{red, green block}}} X_{ij}.$$

The design of this simulation was partially motivated by considering the idea of canalization, where systems operate within appropriate parameters until sufficient perturbations accumulate (e.g., Gibson, 2009). In this third simulation, we set $\beta_j \sim \text{Unif}[0.9, 1.1]$, $\beta_E = 2$ and $\alpha_Q = 1$. We assume that the data have been appropriately normalized, and that the correlation between any two features is greater than or equal to 0. In Simulation 3, we tried to capture the idea that an exposure could lead to coregulation or dysregulation of a cluster of X s, which in itself directly impacts Y . Hence, we defined coregulation as the X s being similar in magnitude and dysregulation, as the X s being very different. The Q_i term in (10) is defined such that the higher values would correspond to strong coregulation, whereas the lower values correspond to dysregulation. For example, suppose Q_i ranges from -5 to 0 . It will be -5 when there are lots of variability (dysregulation), and

0 when there is none (strong coregulation). The function $f(\cdot)$ in (11) simply maps Q_i to the $[0, 1]$ range. In order to get an idea of the relationship in (9), Figure 4 displays the response Y as a function of the first principal component of $\sum_j \beta_j X_{ij}$ (denoted by 1st PC) and $f(Q_i)$. We see that the lower values of $f(Q_i)$ (which implies dysregulation of the features) lead to a lower Y . In this setting, although the clusters do not explicitly include interactions between the X variables, the MARS algorithm allows the possibility of two-way interactions between any of the variables.

3.2.4 | Simulations 4–6

We used the same simulation setup as above, except that we took the continuous outcome Y , defined $p = 1 / (1 + \exp(-Y))$ and used this to generate a two-class outcome z with $\Pr(z = 1) = p$ and $\Pr(z = 0) = 1 - p$. The true parameters were simulated as $\beta_j \sim \text{Unif}[\log(0.9), \log(1.1)]$, $\beta_E = \log(2)$, $\alpha_j \sim \text{Unif}[\log(0.4), \log(0.6)]$ or $\alpha_j \sim \text{Unif}[\log(1.9), \log(2.1)]$. Simulations 4–6 are the binary response versions of simulations 1–3, respectively. The larger odds ratio for E compared to the odds ratio for X is motivated by certain environmental factors that are well known to have substantial impacts on disease risks and phenotypes. For example, body mass index (BMI) has been estimated to explain a large proportion of variation in bone mineral density (BMD) in women (10–20%; Felson, Zhang, Hannan, & Anderson, 1993). This can be converted to a slope of 0.31–0.44 assuming variables are standardized, that is, changes of 0.3–0.4 standard deviations in BMD per standard deviation change in weight. In contrast, the majority of single nucleotide polymorphisms (SNPs) and rare variants have effect sizes under 0.10 standard deviations on BMD (Kemp et al., 2017).

3.3 | Measures of performance

Simulation performance was assessed with measures of model fit, prediction accuracy, and feature stability. Several

TABLE 3 Measures of Performance

Measure	Formula
Model fit	
True-positive rate (TPR)	$ \hat{S} \in S_0 / S_0 $
False-positive rate (FPR)	$ \hat{S} \notin S_0 / j \notin S_0 $
Correct sparsity (Witten, Shojaie, & Zhang, 2014)	$\frac{1}{p} \sum_{j=1}^p A_j$ $A_j = \begin{cases} 1 & \text{if } \hat{\beta}_j = \beta_j = 0 \\ 1 & \text{if } \hat{\beta}_j \neq 0, \beta_j \neq 0 \\ 0 & \text{if else} \end{cases}$
Prediction accuracy	
Root mean squared error (RMSE)	$\ \mathbf{Y}_{test} - \hat{\boldsymbol{\mu}}(\mathbf{X}_{test})\ _2$
Area under the curve (AUC)	Trapezoidal rule
Hosmer–Lemeshow test ($G = 10$)	χ^2 test statistic
Feature stability using K -fold cross-validation on training set (Kalousis, Prados, & Hilario, 2007)	
Pearson Correlation (ρ) (Pearson, 1895)	$\binom{K}{2}^{-1} \sum_{i,j \in \{1, \dots, K\}, i \neq j} \frac{\text{cov}(\hat{\beta}_{(i)}, \hat{\beta}_{(j)})}{\sigma_{\hat{\beta}_{(i)}} \sigma_{\hat{\beta}_{(j)}}}$
Spearman Correlation (r) (Spearman, 1904)	$\binom{K}{2}^{-1} \sum_{i,j \in \{1, \dots, K\}, i \neq j} \left[1 - 6 \sum_m \frac{(r_{m(i)} - r_{m(j)})^2}{p(p^2 - 1)} \right]$
Jaccard Distance (Jaccard, 1912)	$\frac{ \hat{S}_{(i)} \cap \hat{S}_{(j)} }{ \hat{S}_{(i)} \cup \hat{S}_{(j)} }$

^a $\hat{\boldsymbol{\mu}}$: fitting procedure on the training set.

^b S_0 : index of active set = $\{j; \beta_j^0 \neq 0\}$.

^c \hat{S} : index of the set of nonzero estimated coefficients = $\{j; \hat{\beta}_j \neq 0\}$.

^d $|A|$: is the cardinality of set A .

measures for each of these categories, and the specific formulae used are provided in Table 3. We simulated both a training data set and a test data set for each simulation: all tuning parameters for model selection were selected using the training sets only. Although most of the measures of model fit were calculated on the test data sets, true-positive rate, false-positive rate, and correct sparsity were calculated on the training set only. The root mean squared error is determined by predicting the response for the test set using the fitted model on the training set. The area under the curve is determined using the trapezoidal rule (Robin et al., 2011). The stability of feature importance is defined as the variability of feature weights under perturbations of the training set, that is, small modifications in the training set should not lead to considerable changes in the set of important covariates (Tološi & Lengauer, 2011). A feature selection algorithm produces a weight (e.g., $\boldsymbol{\beta} = (\beta_1, \dots, \beta_p)$), a ranking (e.g., $\text{rank}(\boldsymbol{\beta}) : \mathbf{r} = (r_1, \dots, r_m)$) and a subset of features (e.g., $\mathbf{s} = (s_1, \dots, s_p)$, $s_j = \mathbb{I}\{\beta_j \neq 0\}$, where $\mathbb{I}\{\cdot\}$ is the indicator function). In the CLUST and ECLUST methods, we defined a predictor to be nonzero if its corresponding cluster representative weight was nonzero. Using 10-fold cross-validation (CV), we evaluated the similarity between two features and their rankings using Pearson and Spearman correlation, respectively. For each CV fold we reran the models and took the average Pear-

son/Spearman correlations of the $\binom{10}{2}$ combinations of estimated coefficients vectors. To measure the similarity between two subsets of features, we took the average of the Jaccard distance in each fold. A Jaccard distance of 1 indicates perfect agreement between two sets while no agreement will result in a distance of 0. For MARS models we do not report the Pearson/Spearman stability rankings due to the adaptive and functional nature of the model (there are many possible combinations of predictors, each of which are linear basis functions).

3.4 | Results

All reported results are based on 200 simulation runs. We graphically summarized the results across simulations 1–3 for model fit (Figure 5) and feature stability (Figure 6). The results for simulations 4–6 are shown in supplementary Section B, Figures S1–S6. We restrict our attention to $SNR = 1$, $\rho = 0.9$, and $\alpha_j \sim \text{Unif}[1.9, 2.1]$. The model names are labeled as summary measure_model (e.g., avg_lasso corresponds using the average of the features in a cluster as inputs into a lasso regression model). When there is no summary measure appearing in the model name that indicates that the original variables were used (e.g., enet means all separate features were used in the elasticnet model). In Figure 5, panel A, we plot the true-positive rate against

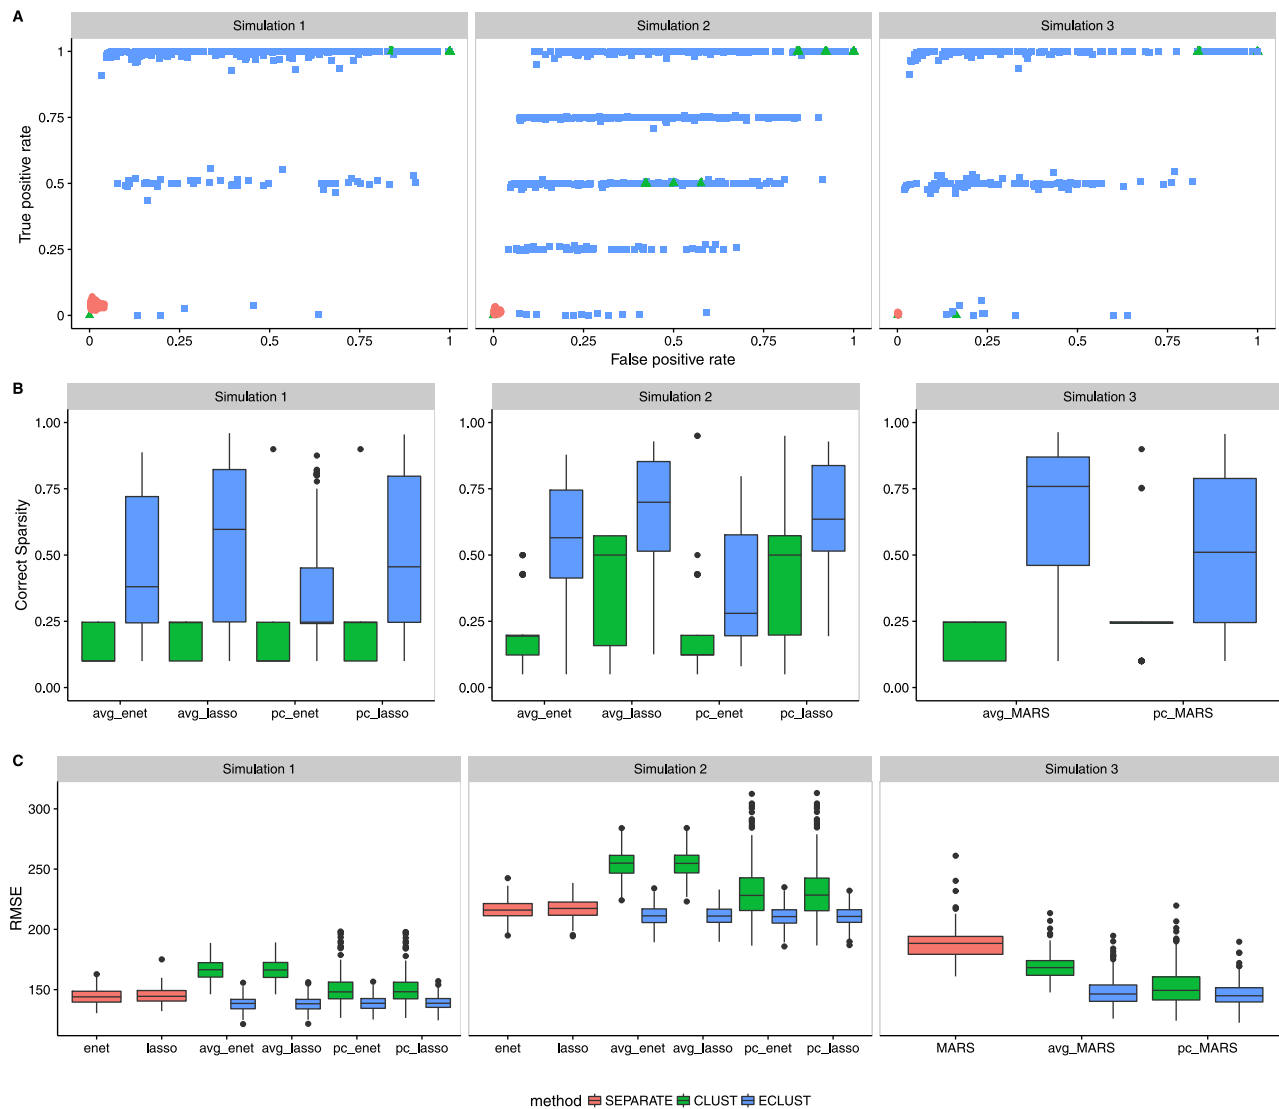


FIGURE 5 Model fit results from simulations 1–3 with $SNR = 1$, $\rho = 0.9$, and $\alpha_j \sim \text{Unif}[1.9, 2.1]$

Note: SEPARATE results are in pink, CLUST in green and ECLUST in blue.

the false-positive rate for each of the 200 simulations. We see that across all simulation scenarios, the SEPARATE method has extremely poor sensitivity compared to both CLUST and ECLUST, which do much better at identifying the active variables, though the resulting models are not always sparse. The relatively few number of green points in panel A is due to the small number of estimated clusters (supplementary Section C, Figure S7) leading to very little variability in performance across simulations. The better performance of ECLUST over CLUST is noticeable as more points lie in the top left part of the plot. The horizontal banding in panel A reflects the stability of the TOM-based clustering approach. ECLUST also does better than CLUST in correctly determining whether a feature is zero or nonzero (Figure 5, panel B). Importantly, across all three simulation scenarios, ECLUST outperforms the competing methods in terms of RMSE (Figure 5, panel

C), regardless of the summary measure and modeling procedure. We present the distribution for the effective number of variables selected in the supplementary material (Figures S8 and S9). We see that the median number of variables selected from ECLUST is less than the median number of variables selected from CLUST, though ECLUST has more variability.

Although the approach using all separate original variables (SEPARATE) produce sparse models, they are sensitive to small perturbations of the data across all stability measures (Figure 6), that is, similar data sets produce very different models. Although the median for the CLUST approach is always slightly better than the median for ECLUST across all stability measures, CLUST results can be much more variable, particularly when stability is measured by the agreement between the value and the ranking of the estimated coefficients across CV folds (Figure 6, panels B and C). The number of

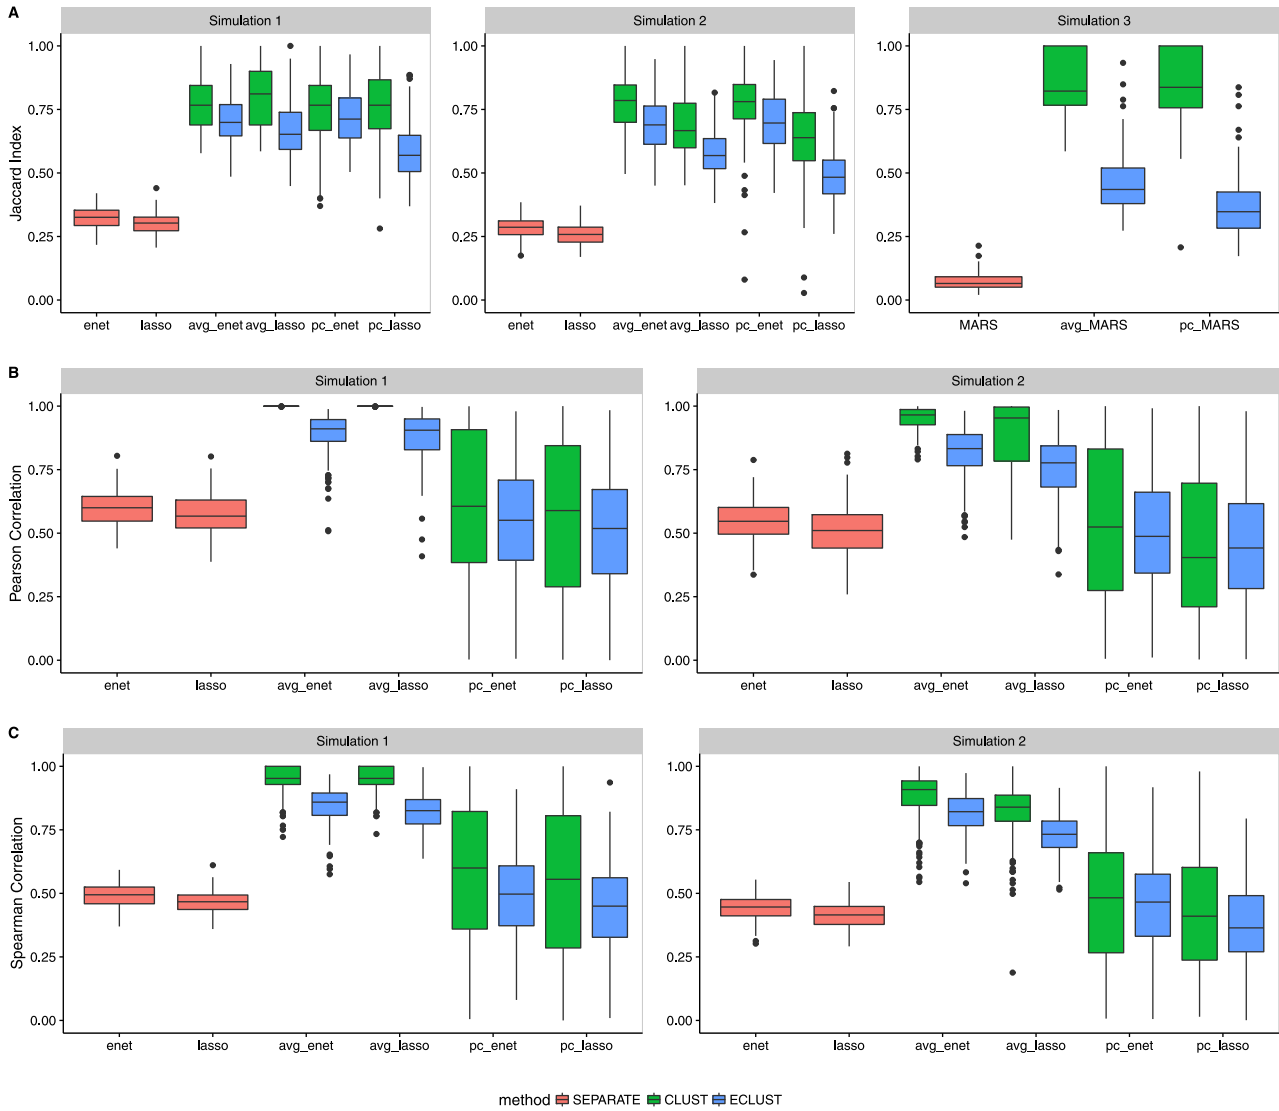


FIGURE 6 Stability results from simulations 1–3 for $SNR = 1$, $\rho = 0.9$, and $\alpha_j \sim \text{Unif}[1.9, 2.1]$

Note: SEPARATE results are in pink, CLUST in green and ECLUST in blue.

estimated clusters, and therefore the number of features in the regression model, tends to be much smaller in CLUST compared to ECLUST, and this explains its poorer performance using the stability measures in Figure 6, because there are more coefficients to estimate. Overall, we observe that the relative performance of ECLUST versus CLUST in terms of stability is consistent across the two summary measures (average or principal component) and across the penalization procedures. The complete results for different values of ρ , SNR and α_j (when applicable) are available in the supplementary Section D, Figures S10–S15 for Simulation 1, Figures S16–S21 for Simulation 2, and Figures S22–S25 for Simulation 3. They show that these conclusions are not sensitive to the SNR , ρ , or α_j . Similar conclusions are made for a binary outcome using logistic regression versions of the lasso, elasticnet, and MARS. ECLUST and CLUST also have better calibration than the SEPARATE method for both linear and nonlinear

models (supplementary Section B, Figures S3–S6). The distributions of Hosmer–Lemeshow (HL) P -values do not follow uniformity. This is in part due to the fact that the HL test has low power in the presence of continuous-dichotomous variable interactions (Hosmer, Hosmer, Le Cessie, & Lemeshow, 1997). Upon inspection of the Q–Q plots, we see that the models have difficulty predicting risks at the boundaries which is a known issue in most models. We also have a small sample size of 200, which means there are on average only 20 subjects in each of the 10 bins. Furthermore, the HL test is sensitive to the choice of bins and method of computing quantiles. Nevertheless, the improved fit relative to the SEPARATE analysis is quite clear.

We also ran all our simulations using the Pearson correlation matrix as a measure of similarity in order to compare its performance against the TOM. The complete results are in supplementary Section E, Figures S26–S31 for Simulation 1,

Figures S32–S37 for Simulation 2, and Figures S38–S41 for Simulation 3. In general, we see slightly better performance of CLUST over ECLUST when using Pearson correlations. This result is probably due to the imprecision in the estimated correlations. The exposure-dependent similarity matrices are quite noisy, and the variability is even larger when we examine the differences between two correlation matrices. Such large levels of variability have a negative impact on the clustering algorithm's ability to detecting the true clusters.

4 | ANALYSIS OF THREE DATA SETS

In this section we demonstrate the performance of ECLUST on three HD data sets with contrasting motivations and features. In the first data set, normal brain development is examined in conjunction with intelligence scores. In the second data set, we aim to identify molecular subtypes of ovarian cancer using gene expression data. The investigators' goal in the third data set is to examine the impact of gestational diabetes mellitus (GDM) on childhood obesity in a sample of mother–child pairs from a prospective birth cohort. The data sets comprise a range of sample sizes, and both amount of clustering in the HD data and strength of the effects of the designated exposure variables vary substantially. Due to the complex nature of these data sets, we decided to use MARS models for Step 2 of our algorithm for all three data sets, as outlined in Table 1. In order to assess performance in these data sets, we have computed the 0.632 estimator (Efron, 1983) and the 95% confidence interval of the R^2 and RMSE from 100 bootstrap samples. The R^2 reported here is defined as the squared Pearson correlation coefficient between the observed and predicted response (Kvålseth, 1985), and the RMSE is defined as in Table 3. Because MARS models can result in unstable predictors (Kuhn, 2008), we also report the results of bagged MARS from $B = 50$ bootstrap samples, where bagging (Breiman, 1996) refers to averaging the predictions from each of the MARS models fit on the B bootstrap samples.

4.1 | NIH MRI study of normal brain development

The NIH MRI Study of Normal Brain Development, started in 2001, was a 7-year longitudinal multisite project that used magnetic resonance technologies to characterize brain maturation in 433 medically healthy, psychiatrically normal children aged 4.5–18 (Evans, A. C., & Brain Development Cooperative Group, 2006). The goal of this study was to provide researchers with a representative and reliable source of healthy control subject data as a basis for understanding atypical brain development associated with a variety of developmental, neurological, and neuropsychiatric disorders

affecting children and adults. Brain imaging data (e.g., cortical surface thickness, intracranial volume), behavioral measures (e.g., IQ scores, psychiatric interviews, behavioral ratings), and demographics (e.g., socioeconomic status) were collected at two year intervals for three time points and are publicly available upon request. Previous research using these data found that level of intelligence and age correlate with cortical thickness (Khundrakpam et al., 2013; Shaw et al., 2006), but to our knowledge no such relation between income and cortical thickness has been observed. We therefore used these data to see the performance of ECLUST in the presence (age) and absence (income) of an effect on the correlations in the HD data. We analyzed the 10,000 most variable regions on the cortical surface from brain scans corresponding to the first sampled time point only. We used binary age ($166 \leq 11.3$ and $172 > 11.3$) and binary income (142 high and 133 low income) indicator as the environment variables and standardized IQ scores as the response. We identified 22 clusters from $TOM(X_{all})$ and 57 clusters from $TOM(X_{diff})$ when using age as the environment, and 86 clusters from $TOM(X_{all})$ and 49 clusters from $TOM(X_{diff})$ when using income as the environment. Results are shown in Figure 7, panels C and D. The method that uses all individual variables as predictors (pink) has better R^2 but also worse RMSE compared to CLUST and ECLUST, likely due to overfitting. There is a slight benefit in performance for ECLUST over CLUST when using age as the environment (panel D). Importantly, we observe very similar performance between CLUST and ECLUST across all models (panel C), suggesting very little impact on the prediction performance when including features derived both with and without the E variable, in a situation where they are unlikely to be relevant.

4.2 | Gene expression study of ovarian cancer

Differences in gene expression profiles have led to the identification of robust molecular subtypes of ovarian cancer; these are of biological and clinical importance because they have been shown to correlate with overall survival (Tothill et al., 2008). Improving prediction of survival time based on gene expression signatures can lead to targeted therapeutic interventions (Helland et al., 2011). The proposed ECLUST algorithm was applied to gene expression data from 511 ovarian cancer patients profiled by the Affymetrix Human Genome U133A 2.0 Array. The data were obtained from the TCGA Research Network: <http://cancergenome.nih.gov/> and downloaded via the TCGA2STAT R library (Wan, Allen, Anderson, & Liu, 2015). Using the 881 signature genes from Helland et al. (2011), we grouped subjects into two groups based on the results in this paper, to create a “positive control” environmental variable expected to have a strong effect. Specifically, we defined an environment variable in our framework as: $E = 0$ for subtypes C1 and

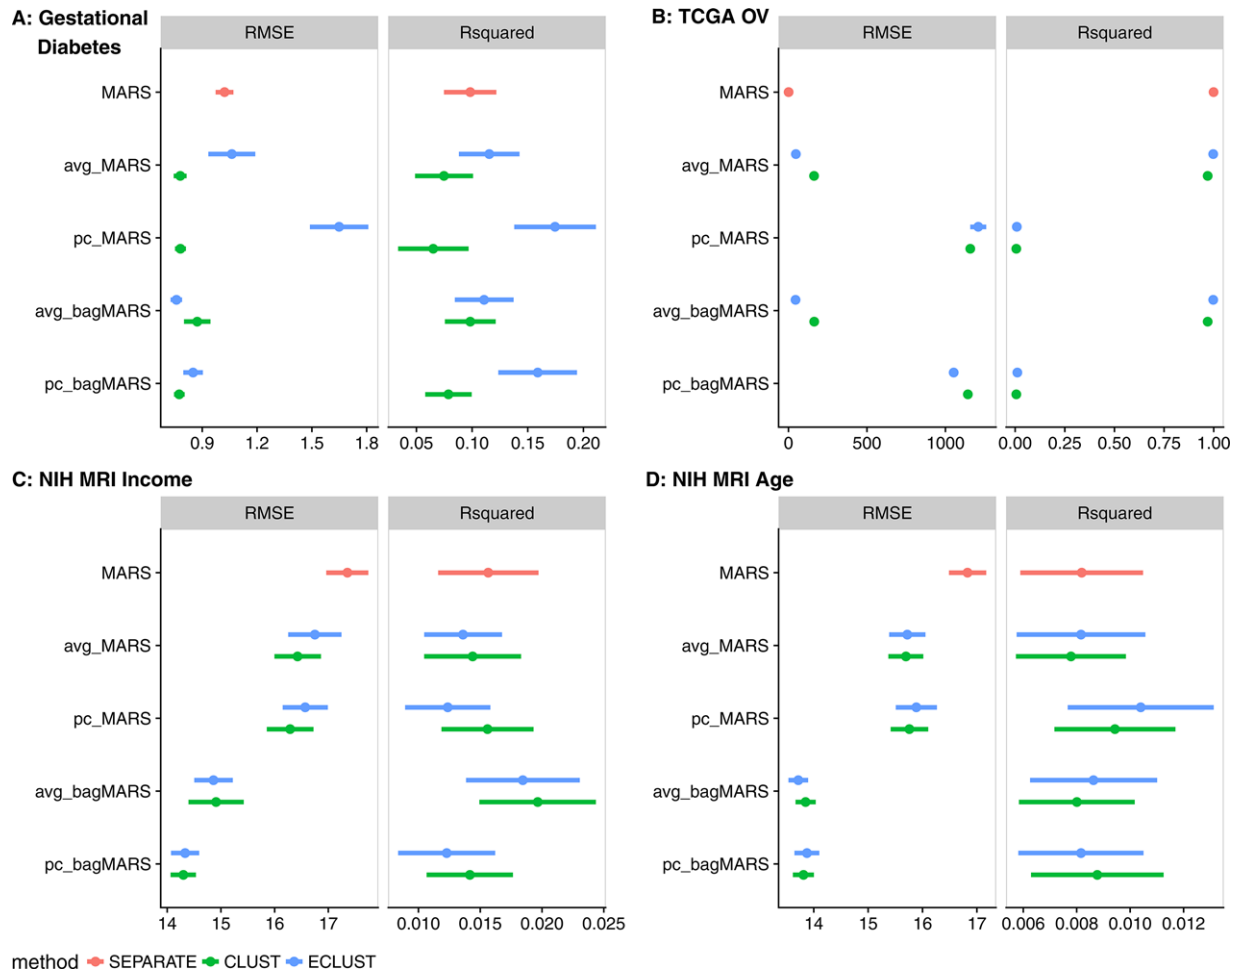


FIGURE 7 Model fit measures from analysis of three data sets: (A) gestational diabetes birth-cohort, (B) TCGA Ovarian Cancer study, (C) NIH MRI Study with income as the environment variable, and (D) NIH MRI Study with age as the environment variable

C2 ($n = 253$), and $E = 1$ for subtypes C4 and C5 ($n = 258$). Overall survival time (log-transformed) was used as the response variable. Because these genes were ascertained on survival time, we expected the method using all genes without clustering to have the best performance, and hence one goal of this analysis was to see if ECLUST performed significantly worse as a result of summarizing the data into a lower dimension. We found three clusters from $TOM(X_{all})$ and three clusters from $TOM(X_{diff})$; results are shown in Figure 7, panel C. Across all models, ECLUST performs slightly better than CLUST. Furthermore, it performs almost as well as the separate variable method, with the added advantage of dealing with a much smaller number of predictors (881 with SEPARATE compared to 6 with ECLUST).

4.3 | Gestational diabetes, epigenetics, and metabolic disease

Events during pregnancy are suspected to play a role in childhood obesity development but only little is known about the mechanisms involved. Indeed, children born to women

who had GDM in pregnancy are more likely to be overweight and obese (Wendland et al., 2012), and evidence suggests epigenetic factors are important piece of the puzzle (Bouchard et al., 2010, 2012). Recently, methylation changes in placenta and cord blood were associated with GDM (Ruchat et al., 2013), and here we explore how these changes are associated with obesity in the children at the age of about 5. DNA methylation in placenta was measured with the Infinium HumanMethylation450 BeadChip (Illumina, Inc.; Bibikova et al., 2011) microarray in a sample of 28 women, 20 of whom had a GDM-affected pregnancy, and here we used GDM status as our E variable, assuming that this has widespread effects on DNA methylation and on its correlation patterns. Our response, Y , is the standardized BMI in the offspring at the age of 5. In contrast to the previous two examples, here we had no particular expectation of how ECLUST would perform. Using the 10,000 most variable probes, we found 2 clusters from $TOM(X_{all})$ and 75 clusters from $TOM(X_{diff})$. The predictive model results from a MARS analysis are shown in Figure 7, panel A. When using R^2 as the measure of performance, ECLUST outperforms both SEPARATE

TABLE 4 Ingenuity Pathway Analysis results—top-ranked diseases and disorders, and physiological system development and function epigenetically affected by gestational diabetes mellitus and associated with childhood body mass index

Category	Name	<i>P</i> values	<i>n</i> ^a
Diseases and disorders	Hepatic system disease	[9.61e−7 to 5.17e−7]	75
Physiological system development and function			
	Behavior	[1.35e−2 to 7.82e−8]	33
	Embryonic development	[1.35e−2 to 2.63e−8]	26
	Nervous system development and function	[1.35e−2 to 2.63e−8]	43
	Organ development	[1.35e−2 to 2.63e−8]	20
	Organismal development	[1.35e−2 to 2.63e−8]	34

^aNumber of genes involved in each pathway.

and CLUST methods. When using RMSE as the measure of model performance, performance tended to be better with CLUST rather than ECLUST perhaps in part due to the small number of clusters derived from $TOM(X_{all})$ relative to $TOM(X_{diff})$. Overall, the ECLUST algorithm with bagged MARS and the first PC of each cluster performed best, that is, it had a better R^2 than CLUST with comparable RMSE. The sample size here is very small, and therefore the stability of the model fits is limited. The probes in these clusters mapped to 164 genes and these genes were selected to conduct pathway analyses using the Ingenuity Pathway Analysis (IPA) software (Ingenuity System). IPA compares the selected genes to a reference list of genes included in many biological pathways using a hypergeometric test. Smaller *P* values are evidence for overrepresented gene ontology categories in the input gene list. The results are summarized in Table 4 and provide some biological validation of our ECLUST method. For example, the hepatic system is involved with the metabolism of glucose and lipids (Saltiel & Kahn, 2001), and behavior and neurodevelopment are associated with obesity (Epstein, Paluch, Kilanowski, & Raynor, 2004). Furthermore, it is interesting that embryonic and organ development pathways are involved because GDM is associated with macrosomia (Ehrenberg, Mercer, & Catalano, 2004).

5 | DISCUSSION

The challenge of precision medicine is to appropriately fit treatments or recommendations to each individual. Data such as gene expression, DNA methylation levels, or magnetic resonance imaging (MRI) signals are examples of HD measurements that capture multiple aspects of how a tissue is functioning. These data often show patterns associated with disease, and major investments are being made in the genomics research community to generate such HD data. Analytic tools increasing prediction accuracy are needed to maximize the productivity of these investments. However, the effects of exposures have usually been overlooked, but these are cru-

cial because they can lead to ways to intervene. Hence, it is essential to have a clear understanding of how exposures modify HD measures, and how the combination leads to disease. Existing methods for prediction (of disease), which are based on HD data and interactions with exposures, fall far short of being able to obtain this clear understanding. Most methods have low power and poor interpretability, and furthermore, modeling and interpretation problems are exacerbated when there is interest in interactions. In general, power to estimate interactions is low, and the number of possible interactions could be enormous. Therefore, here we have proposed a strategy to leverage situations where a covariate (e.g., an exposure) has a widespread effect on one or more HD measures, for example, GDM on methylation levels. We have shown that this expected pattern can be used to construct dimension-reduced predictor variables that inherently capture the systemic covariate effects. These dimension-reduced variables, constructed without using the phenotype, can then be used in predictive models of any type. In contrast to some common analysis strategies that model the effects of individual predictors on outcome, our approach makes a step toward a systems-based perspective that we believe will be more informative when exploring the factors that are associated with disease or a phenotype of interest. We have shown, through simulations and real data analysis, that incorporation of environmental factors into predictive models in a way that retains a HD perspective can improve results and interpretation for both linear and nonlinear effects.

We proposed two key methodological steps necessary to maximize predictive model interpretability when using HD data and a binary exposure: (1) dimension reduction of HD data built on exposure sensitivity and (2) implementation of penalized prediction models. In the first step, we proposed to identify exposure-sensitive HD pairs by contrasting the TOM between exposed and unexposed individuals; then we cluster the elements in these HD pairs to find exposure-sensitive coregulated sets. New dimension-reduced variables that capture exposure-sensitive features (e.g., the first principal component of each cluster) were then defined. In the second step

we implemented linear and nonlinear variable selection methods using the dimension-reduced variables to ensure stability of the predictive model. The ECLUST method has been implemented in the `eclust` (Bhatnagar, 2017) R package publicly available on CRAN. Our method along with computationally efficient algorithms, allows the analysis of up to 10,000 variables at a time on a laptop computer.

The methods that we have proposed here are currently only applicable when three data elements are available. Specifically a binary environmental exposure, an HD data set that can be affected by the exposure, and a single phenotype. When comparing the TOM and Pearson correlations as a measure of similarity, our simulations showed that the performance of ECLUST was worse with correlations. This speaks to the potential of developing a better measure than the difference of two matrices. For example, we are currently exploring ways in which to handle continuous exposures or multiple exposures. The best way to construct an exposure-sensitive distance matrix that can be used for clustering is not obvious in these situations. One possible solution relies on a nonparametric smoothing based approach where weighted correlations are calculated. These weights can be derived from a kernel-based summary of the exposure covariates (e.g., Qiu, Han, Liu, & Caffo, 2016). Then, contrasting unweighted and weighted matrices will allow construction of covariate-sensitive clusters. The choice of summary measure for each cluster also warrants further study. Although principal components and averages are well understood and easy to implement, the main shortcoming is that they involve all original variables in the group. As the size of the groups increase, the interpretability of these measures decreases. Nonnegative matrix factorization (Lee & Seung, 2001) and sparse principal component analysis (SPCA; Witten, Tibshirani, & Hastie, 2009) are alternatives that find sparse and potentially interpretable factors. Furthermore, structured SPCA (Jenatton, Obozinski, & Bach, 2010) goes beyond restricting the cardinality of the contributing factors by imposing some a priori structural constraints deemed relevant to model the data at hand.

We are all aware that our exposures and environments impact our health and risks of disease, however detecting how the environment acts is extremely difficult. Furthermore, it is very challenging to develop reliable and understandable ways of predicting the risk of disease in individuals, based on HD data such as genomic or imaging measures, and this challenge is exacerbated when there are environmental exposures that lead to many subtle alterations in the genomic measurements. Hence, we have developed an algorithm and an easy-to-use software package to transform analysis of how environmental exposures impact human health, through an innovative signal-extracting approach for HD measurements. Evidently, the model fitting here is performed using existing methods; our goal is to illustrate the potential of improved dimension reduction in two-stage methods, in order to generate discussion

and new perspectives. If such an approach can lead to more interpretable results that identify gene–environment interactions and their effects on diseases and traits, the resulting understanding of how exposures influence the high-volume measurements now available in precision medicine will have important implications for health management and drug discovery.

AVAILABILITY OF DATA AND MATERIAL

1. NIH MRI Study of Normal Brain Development data are available in the Pediatric MRI Data Repository, <https://pediatricmri.nih.gov/>
2. Gene Expression Study of Ovarian Cancer data are available in the Genomic Data Commons repository, <https://gdc.cancer.gov/>, and were downloaded via the TCGA2STAT R library Wan et al. (2015)
3. Gestational diabetes, epigenetics, and metabolic disease: the clinical data, similarity matrices and cluster summaries are available at Zenodo [10.5281/zenodo.259222]. The raw analyzed during the current study are not publicly available due to reasons of confidentiality, although specific collaborations with LB can be requested.

ACKNOWLEDGMENTS

SRB was supported by the Ludmer Centre for Neuroinformatics and Mental Health.

SRB, CMTG, YY, and MB contributed to the conceptualization of this research; SRB, LB, and BK contributed to the data curation; SRB contributed to the formal analysis, software, visualization; SRB and CMTG contributed to the methodology; SRB and CMTG contributed to writing the original draft; SRB, CMTG, YY, BK, ACE, MB, and LB contributed to writing, reviewing, and editing the draft.

CONFLICT OF INTEREST

The authors declare no conflicts of interest.

ORCID

Sahir Rai Bhatnagar 
<http://orcid.org/0000-0001-8956-2509>
 Celia M.T. Greenwood 
<http://orcid.org/0000-0002-2427-5696>

REFERENCES

- Bhatnagar, S. R. (2017). *eclust: environment based clustering for interpretable predictive models in high dimensional data*. R package version 0.1.0. Retrieved from <https://github.com/sahirbhatnagar/eclust/>

- Bibikova, M., Barnes, B., Tsan, C., Ho, V., Klotzle, B., Le, J. M., ... Shen, R. (2011). High density dna methylation array with single CpG site resolution. *Genomics*, 98(4), 288–295.
- Bouchard, L., Hivert, M.-F., Guay, S.-P., St-Pierre, J., Perron, P., & Brisson, D. (2012). Placental adiponectin gene dna methylation levels are associated with mothers' blood glucose concentration. *Diabetes*, 61(5), 1272–1280.
- Bouchard, L., Thibault, S., Guay, S.-P., Santure, M., Monpetit, A., St-Pierre, J., ... Brisson, D. (2010). Leptin gene epigenetic adaptation to impaired glucose metabolism during pregnancy. *Diabetes Care*, 33(11), 2436–2441.
- Breiman, L. (1996). Bagging predictors. *Machine Learning*, 24(2), 123–140.
- Bühlmann, P., Rütimann, P., van de Geer, S., & Zhang, C.-H. (2013). Correlated variables in regression: clustering and sparse estimation. *Journal of Statistical Planning and Inference*, 143(11), 1835–1858.
- Efron, B. (1983). Estimating the error rate of a prediction rule: improvement on cross-validation. *Journal of the American Statistical Association*, 78(382), 316–331.
- Ehrenberg, H. M., Mercer, B. M., & Catalano, P. M. (2004). The influence of obesity and diabetes on the prevalence of macrosomia. *American Journal of Obstetrics and Gynecology*, 191(3), 964–968.
- Epstein, L. H., Paluch, R. A., Kilanowski, C. K., & Raynor, H. A. (2004). The effect of reinforcement or stimulus control to reduce sedentary behavior in the treatment of pediatric obesity. *Health Psychology*, 23(4), 371.
- Evans, A. C., & Brain Development Cooperative Group. (2006). The NIH MRI study of normal brain development. *Neuroimage*, 30(1), 184–202.
- Fan, J., Han, F., & Liu, H. (2014). Challenges of big data analysis. *National Science Review*, 1(2), 293–314.
- Felson, D. T., Zhang, Y., Hannan, M. T., & Anderson, J. J. (1993). Effects of weight and body mass index on bone mineral density in men and women: the Framingham study. *Journal of Bone and Mineral Research*, 8(5), 567–573.
- Friedman, J., Hastie, T., & Tibshirani, R. (2010). Regularization paths for generalized linear models via coordinate descent. *Journal of Statistical Software*, 33(1), 1.
- Friedman, J. H. (1991). Multivariate adaptive regression splines. *Annals of Statistics*, 1–67.
- Gibson, G. (2009). Decanalization and the origin of complex disease. *Nature Reviews Genetics*, 10(2), 134.
- Hastie, T., Tibshirani, R., Botstein, D., & Brown, P. (2001). Supervised harvesting of expression trees. *Genome Biology*, 2(1), 1–0003.
- Helland, Å., Anglesio, M. S., George, J., Cowin, P. A., Johnstone, C. N., House, C. M., ... Rustgi, A. K. (2011). Deregulation of MYCN, LIN28b and LET7 in a molecular subtype of aggressive high-grade serous ovarian cancers. *PLoS One*, 6(4), e18064.
- Hosmer, D. W., Hosmer, T., Le Cessie, S., & Lemeshow, S. (1997). A comparison of goodness-of-fit tests for the logistic regression model. *Statistics in Medicine*, 16(9), 965–980.
- Jaccard, P. (1912). The distribution of the flora in the alpine zone. *New Phytologist*, 11(2), 37–50.
- Jenatton, R., Obozinski, G., & Bach, F. (2010). Structured sparse principal component analysis. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics* (pp. 366–373).
- Kalousis, A., Prados, J., & Hilario, M. (2007). Stability of feature selection algorithms: a study on high-dimensional spaces. *Knowledge and Information Systems*, 12(1), 95–116.
- Kemp, J. P., Morris, J. A., Medina-Gomez, C., Forgetta, V., Warrington, N. M., Youlten, S. E., ... Evans, D. M. (2017). Identification of 153 new loci associated with heel bone mineral density and functional involvement of gpc6 in osteoporosis. *Nature Genetics*, 49, 1468.
- Kendall, M. (1957). *A course in multivariate analysis*. London: Griffin.
- Khundrakpam, B. S., Reid, A., Brauer, J., Carbonell, F., Lewis, J., Ameis, S., ... Brain Development Cooperative Group. (2013). Developmental changes in organization of structural brain networks. *Cerebral Cortex*, 23(9), 2072–2085.
- Kostka, D., & Spang, R. (2004). Finding disease specific alterations in the co-expression of genes. *Bioinformatics*, 20(Suppl. 1), i194–i199.
- Kuhn, M. (2008). caret package. *Journal of Statistical Software*, 28(5). Retrieved from <http://www.jstatsoft.org/index>
- Kvålseth, T. O. (1985). Cautionary note about r^2 . *American Statistician*, 39(4), 279–285.
- Langfelder, P., & Horvath, S. (2007). Eigengene networks for studying the relationships between co-expression modules. *BMC Systems Biology*, 1(1), 54.
- Langfelder, P., & Horvath, S. (2008). Wgcna: an R package for weighted correlation network analysis. *BMC Bioinformatics*, 9(1), 1.
- Langfelder, P., Zhang, B., & Horvath, S. (2008). Defining clusters from a hierarchical cluster tree: the dynamic tree cut package for R. *Bioinformatics*, 24(5), 719–720.
- Langfelder, P., Zhang, B., & with contributions from Horvath, S. (2016). *dynamicTreeCut: methods for detection of clusters in hierarchical clustering dendrograms*. R package version 1.63-1. Retrieved from <http://www.genetics.ucla.edu/labs/horvath/CoexpressionNetwork/BranchCutting/>
- Lee, D. D., & Seung, H. S. (2001). Algorithms for non-negative matrix factorization. In T. K. Leen, T. G. Dietterich, & V. Tresp (Eds.), *Advances in neural information processing systems* (Vol. 13, pp. 556–562). Cambridge, MA: MIT Press.
- Milborrow, S. (2011). Derived from mda: MARS by Trevor Hastie and Rob Tibshirani. *Earth: multivariate adaptive regression splines*. R package version 3.2-0. Retrieved from <http://CRAN.R-project.org/package=earth>.
- Müllner, D. (2013). fastcluster: fast hierarchical, agglomerative clustering routines for R and Python. *Journal of Statistical Software*, 53(9), 1–18.
- Oros Klein, K., Oualkacha, K., Lafond, M.-H., Bhatnagar, S., Tonin, P. N., & Greenwood, C. M. (2016). Gene coexpression analyses differentiate networks associated with diverse cancers harboring tp53 missense or null mutations. *Frontiers in Genetics*, 7, 137.
- Park, M. Y., Hastie, T., & Tibshirani, R. (2007). Averaged gene expressions for regression. *Biostatistics*, 8(2), 212–227.
- Pearson, K. (1895). Note on regression and inheritance in the case of two parents. *Proceedings of the Royal Society of London*, 240–242.
- Qiu, H., Han, F., Liu, H., & Caffo, B. (2016). Joint estimation of multiple graphical models from high dimensional time series. *Journal of the*

- Royal Statistical Society: *Series B (Statistical Methodology)*, 78(2), 487–504.
- R Core Team. (2016). *R: a language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing.
- Ravasz, E., Somera, A. L., Mongru, D. A., Oltvai, Z. N., & Barabási, A.-L. (2002). Hierarchical organization of modularity in metabolic networks. *Science*, 297(5586), 1551–1555.
- Reid, S., & Tibshirani, R. (2016). Sparse regression and marginal testing using cluster prototypes. *Biostatistics*, 17(2), 364–376.
- Robin, X., Turck, N., Hainard, A., Tiberti, N., Lisacek, F., Sanchez, J.-C., & MÄijller, M. (2011). proc: an open-source package for r and s+ to analyze and compare roc curves. *BMC Bioinformatics*, 12, 77.
- Ruchat, S.-M., Houde, A.-A., Voisin, G., St-Pierre, J., Perron, P., Bailargeon, J.-P., ... Bouchard, L. (2013). Gestational diabetes mellitus epigenetically affects genes predominantly involved in metabolic diseases. *Epigenetics*, 8(9), 935–943.
- Saltiel, A. R., & Kahn, C. R. (2001). Insulin signalling and the regulation of glucose and lipid metabolism. *Nature*, 414(6865), 799–806.
- Sathirapongsasuti, J. F. (2013). *COPDSexualDimorphism.data: data to support sexually dimorphic and COPD differential analysis for gene expression and methylation*. R package version 1.4.0. Retrieved from <https://git.bioconductor.org/packages/COPDSexualDimorphism.data>
- Sato, J. R., Hoexter, M. Q., de Magalhães Oliveira, P. P., Brammer, M. J., Murphy, D., ... Ecker, C. (2013). Inter-regional cortical thickness correlations are associated with autistic symptoms: a machine-learning approach. *Journal of Psychiatric Research*, 47(4), 453–459.
- Shaw, P., Greenstein, D., Lerch, J., Clasen, L., Lenroot, R., Gogtay, N. e. a., ... Giedd, J. (2006). Intellectual ability and cortical development in children and adolescents. *Nature*, 440(7084), 676–679.
- Spearman, C. (1904). The proof and measurement of association between two things. *American Journal of Psychology*, 15(1), 72–101.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, 58(1), 267–288.
- Tološi, L., & Lengauer, T. (2011). Classification with correlated features: unreliability of feature ranking and solutions. *Bioinformatics*, 27(14), 1986–1994.
- Tothill, R. W., Tinker, A. V., George, J., Brown, R., Fox, S. B., Lade, S., ... Bowtell, D. D. (2008). Novel molecular subtypes of serous and endometrioid ovarian cancer linked to clinical outcome. *Clinical Cancer Research*, 14(16), 5198–5208.
- Wan, Y.-W., Allen, G. I., Anderson, M. L., & Liu, Z. (2015). *TCGA2STAT: simple TCGA data access for integrated statistical analysis in R*. R package version 1.2. Retrieved from <http://www.liuzlab.org/TCGA2STAT/>
- Wendland, E. M., Torloni, M. R., Falavigna, M., Trujillo, J., Dode, M. A., Campos, M. A., ... Schmidt, M. I. (2012). Gestational diabetes and pregnancy outcomes—a systematic review of the World Health Organization (WHO) and the International Association of Diabetes in Pregnancy Study Groups (IADPSG) diagnostic criteria. *BMC Pregnancy and Childbirth*, 12(1), 1.
- Witten, D. M., Shojai, A., & Zhang, F. (2014). The cluster elastic net for high-dimensional regression with unknown variable grouping. *Technometrics*, 56(1), 112–122.
- Witten, D. M., Tibshirani, R., & Hastie, T. (2009). A penalized matrix decomposition, with applications to sparse principal components and canonical correlation analysis. *Biostatistics*, 10(3), 515–534.
- Yang, Y., & Zou, H. (2014). gglasso: group lasso penalized learning using a unified bmd algorithm. R package version 1.3. Retrieved from <https://github.com/emeryyi/gglasso>
- Zhang, B., & Horvath, S. (2005). A general framework for weighted gene co-expression network analysis. *Statistical Applications in Genetics and Molecular Biology*, 4(1), Epub 2005. <https://doi.org/10.2202/1544-6115.1128>
- Zou, H., & Hastie, T. (2005). Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2), 301–320.

SUPPORTING INFORMATION

Additional Supporting Information may be found online in the supporting information tab for this article.

How to cite this article: Bhatnagar SR, Yang Y, Khundrakpam B, Evans AC, Blanchette M, Bouchard L, Greenwood CMT. An analytic approach for interpretable predictive models in high-dimensional data in the presence of interactions with exposures. *Genet Epidemiol*. 2018; 1–17. <https://doi.org/10.1002/gepi.22112>

Supplemental Material: An analytic approach for interpretable predictive models in high dimensional data, in the presence of interactions with exposures

Sahir Rai Bhatnagar, Yi Yang, Budhachandra Khundrakpam, Alan Evans,
Mathieu Blanchette, Luigi Bouchard, Celia MT Greenwood

October 4, 2017

Contents

A Supplemental Methods	1
A.1 Description of Topological Overlap Matrix	1
B Binary Outcome Simulation Results	2
C Analysis of Clusters	8
D Simulation Results Using TOM as a Measure of Similarity	10
D.1 Simulation 1	10
D.2 Simulation 2	16
D.3 Simulation 3	22
E Simulation Results Using Pearson Correlations as a Measure of Similarity	26
E.1 Simulation 1	26
E.2 Simulation 2	32
E.3 Simulation 3	38
F Visual Representation of Similarity Matrices	42
F.1 Pearson Correlation Matrix	42

A Supplemental Methods

A.1 Description of Topological Overlap Matrix

Starting with a similarity measure $s_{ij} = |cor(i, j)|$ between node i and node j , one could apply a hard threshold to determine if this pair is considered connected or not resulting in an un-weighted network (a matrix of 0's and 1's). Instead, Zhang and Horvath (Zhang and Horvath, 2005) propose a soft thresholding framework that assigns a connection weight to each gene pair using a power adjacency function $a_{ij} = |s_{ij}|^\beta$. The parameter β determines the sensitivity and specificity of the pairwise connection strengths e.g. a larger β will result in fewer connected nodes which can reduce noise in the network but can also eliminate signal if too large. A measure of similarity is then derived using the symmetric and non-negative topological overlap matrix (Ravasz et al., 2002) (TOM) $\Omega = [\omega_{ij}]$:

$$\omega_{ij} = \frac{l_{ij} + a_{ij}}{\min\{k_i, k_j\} + 1 - a_{ij}} \quad (1)$$

where $l_{ij} = \sum_u a_{iu}a_{uj}$, $k_i = \sum_u a_{iu}$ is the node connectivity, and the index u runs across all nodes of the network. Basically, ω_{ij} is a measure of similarity in terms of the commonality of the nodes they connect to. If i and j are unconnected and do not share any neighbors then $\omega_{ij} = 0$. An $\omega_{ij} = 1$ means that i and j are connected, and the neighbors of the node with fewer connections are also neighbors of the other node.

B Binary Outcome Simulation Results

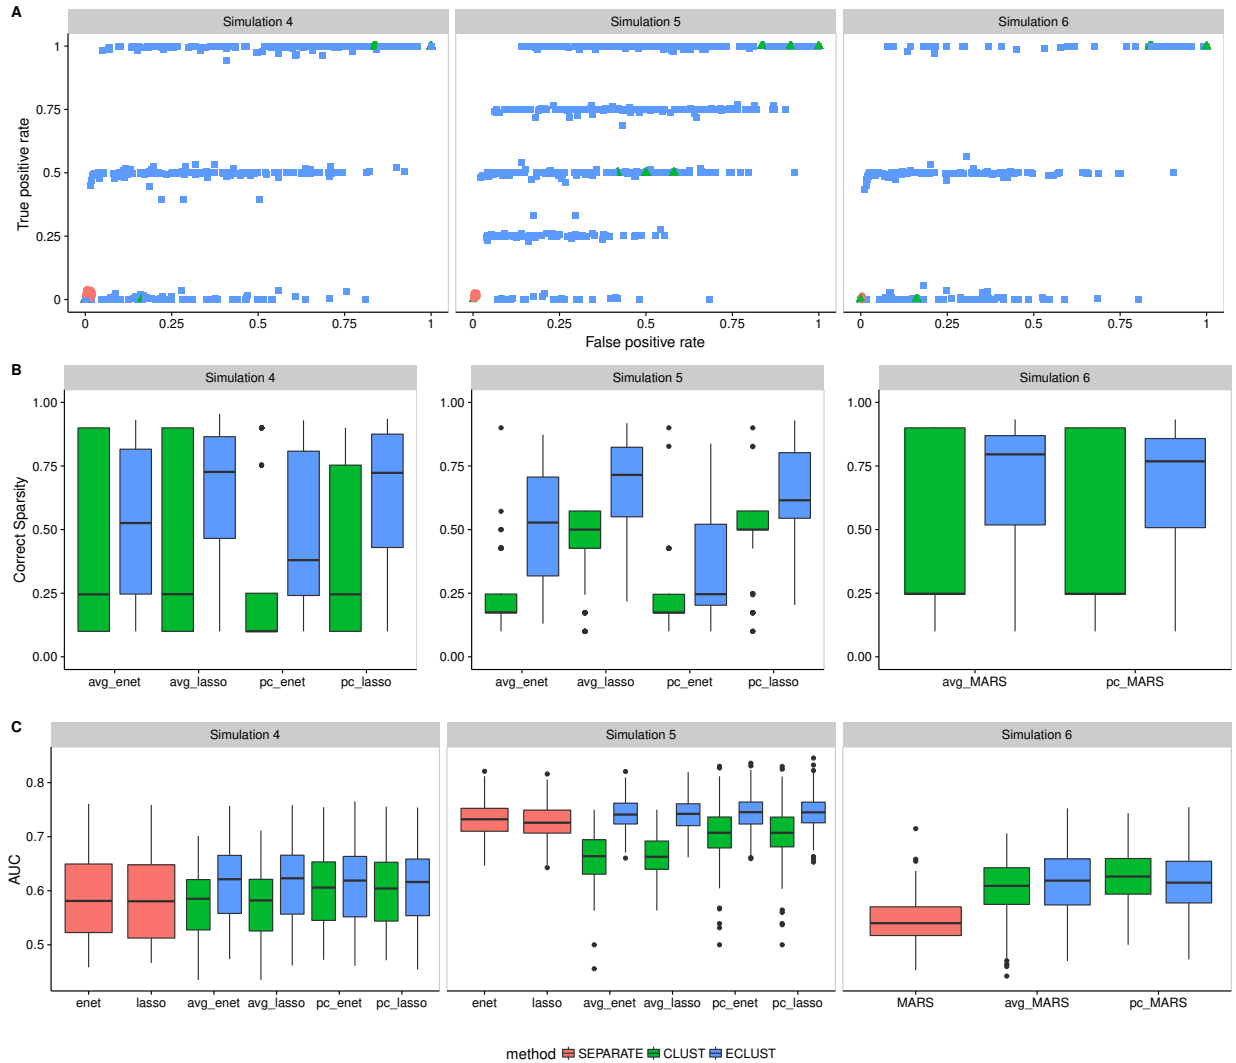


Figure S1: Model fit results from simulations 4, 5 and 6 for $SNR = 1$, $\rho = 0.9$, and $\alpha_j \sim \text{Unif}[\log(1.9), \log(2.1)]$. SEPARATE results are in pink, CLUST in green and ECLUST in blue.

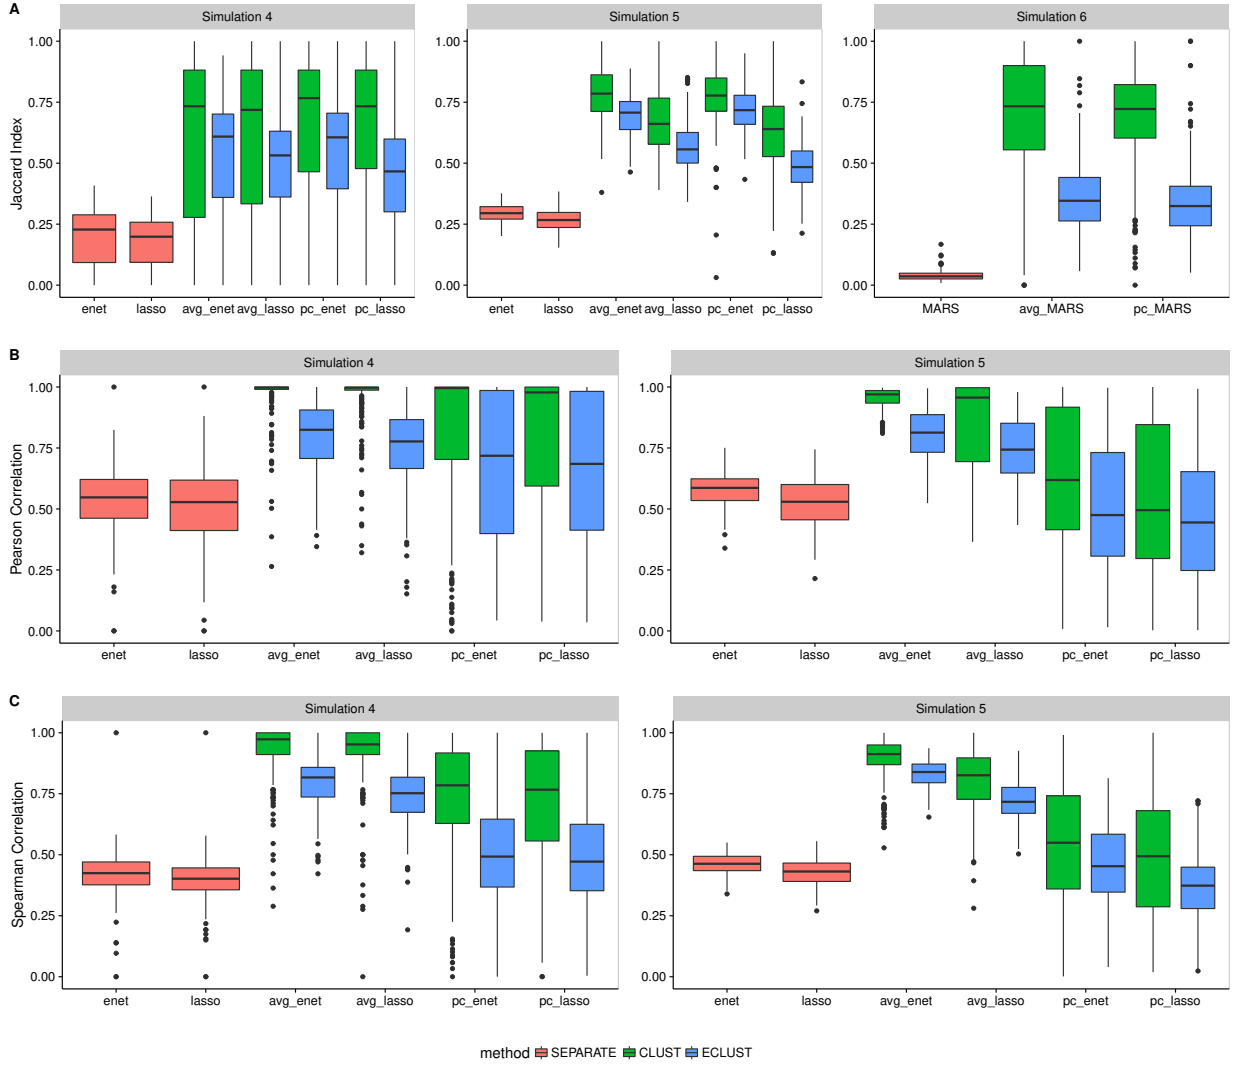


Figure S2: Stability results from simulations 4, 5 and 6 for $SNR = 1$, $\rho = 0.9$, and $\alpha_j \sim \text{Unif}[\log(1.9), \log(2.1)]$. SEPARATE results are in pink, CLUST in green and ECLUST in blue.




Figure S3: Hosmer-Lemeshow statistics from simulations 4, 5 and 6 for $SNR = 1$, $\rho = 0.9$, and $\alpha_j \sim \text{Unif}[\log(1.9), \log(2.1)]$. SEPARATE results are in pink, CLUST in green and ECLUST in blue.

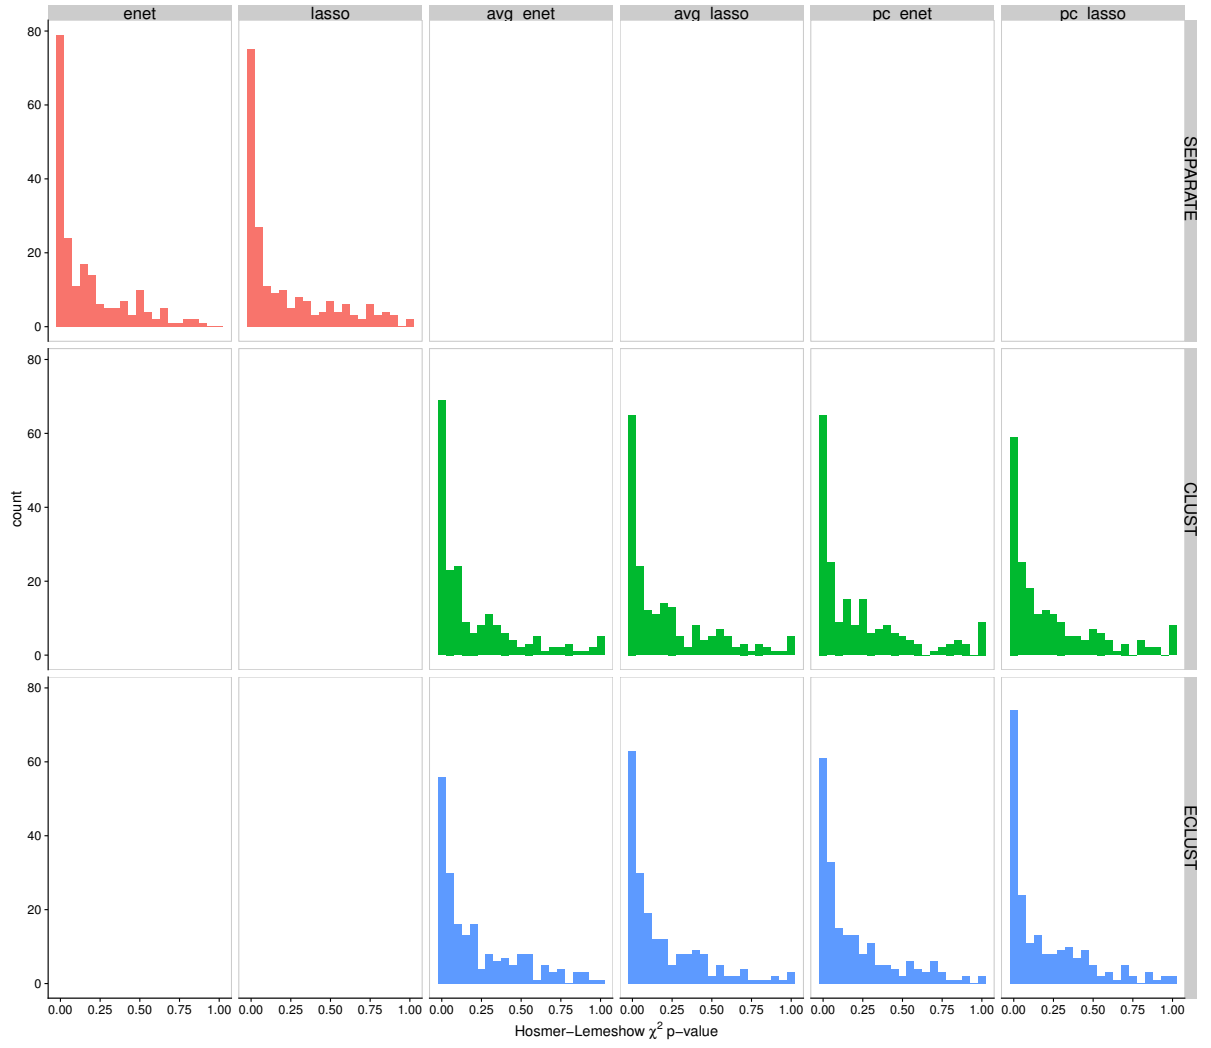


Figure S4: Hosmer-Lemeshow p-values from simulation 4 for $SNR = 1$, $\rho = 0.9$, and $\alpha_j \sim \text{Unif}[\log(1.9), \log(2.1)]$. SEPARATE results are in pink, CLUST in green and ECLUST in blue.

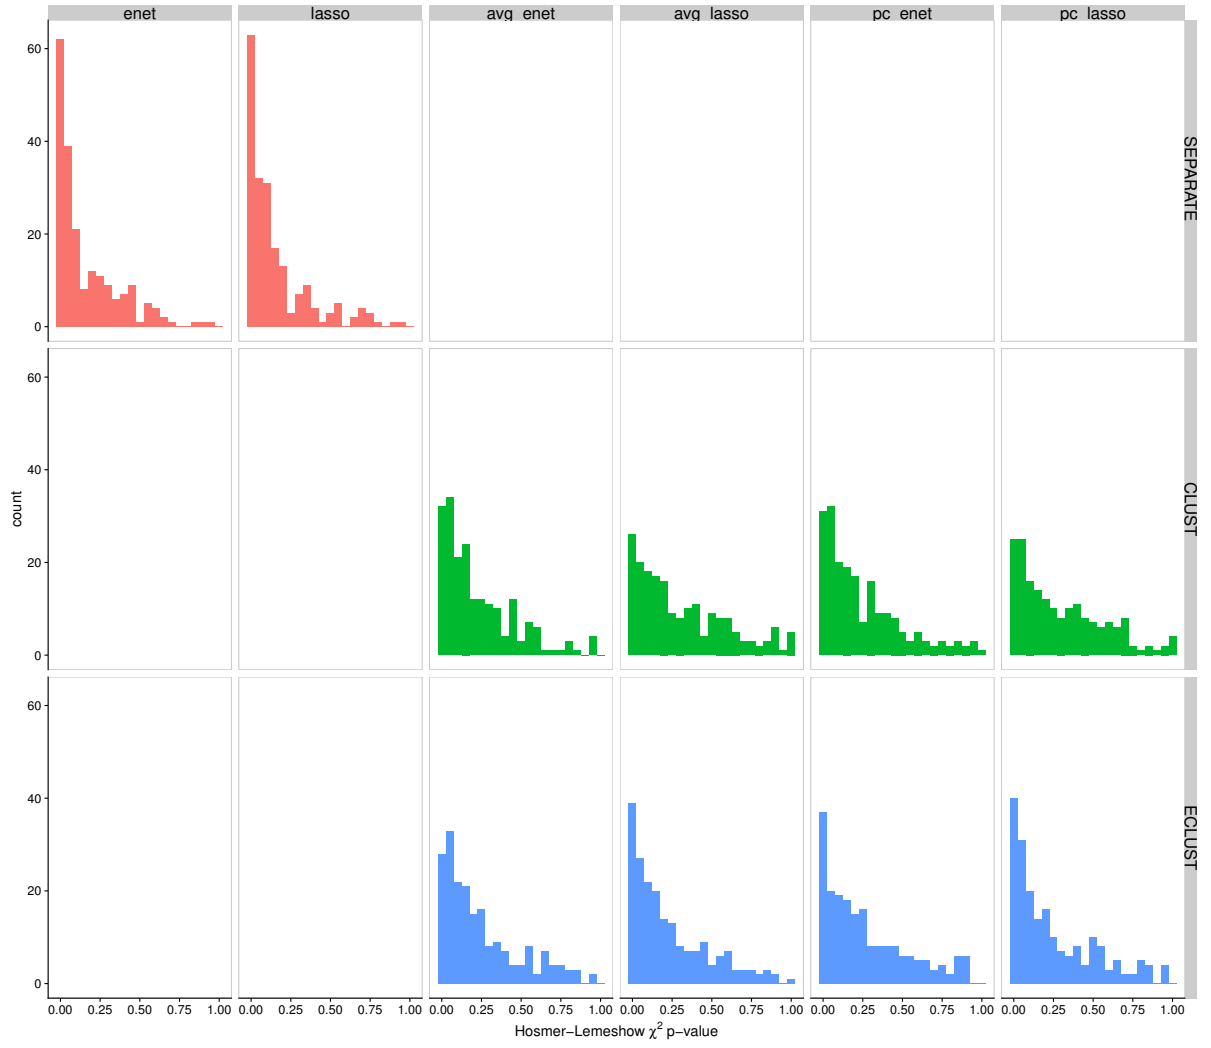


Figure S5: Hosmer-Lemeshow p-values from simulation 5 for $SNR = 1$, $\rho = 0.9$, and $\alpha_j \sim \text{Unif}[\log(1.9), \log(2.1)]$. SEPARATE results are in pink, CLUST in green and ECLUST in blue.




Figure S6: Hosmer-Lemeshow p-values from simulation 6 for $SNR = 1$, $\rho = 0.9$, and $\alpha_j \sim \text{Unif}[\log(1.9), \log(2.1)]$. SEPARATE results are in pink, CLUST in green and ECLUST in blue.

C Analysis of Clusters




Figure S7: Number of estimated clusters from applying the `dynamicTreeCut` algorithm to hierarchical clustering of the dissimilarity matrix with average linkage. Left panel: CLUST uses $1 - \text{Cor}(X_{all})$ and ECLUST uses the euclidean distance of $\text{Cor}(X_{diff})$ as measures of dissimilarity. Right panel: CLUST uses $1 - \text{TOM}(X_{all})$ and ECLUST uses the euclidean distance of $\text{TOM}(X_{diff})$ as measures of dissimilarity. Empirical distributions based on 200 simulation runs.




Figure S8: Effective number of selected variables for simulations 1-3 for $SNR = 1, \rho = 0.9$. A variable was considered “selected” if its corresponding cluster representative was selected.




Figure S9: Effective number of selected variables for simulations 4-6 for $SNR = 1$, $\rho = 0.9$ and $\alpha_j \sim \text{Unif}[\log(1.9), \log(2.1)]$. A variable was considered “selected” if its corresponding cluster representative was selected.

D Simulation Results Using TOM as a Measure of Similarity

D.1 Simulation 1




Figure S10: Simulation 1 – Root mean squared error on an independent test set using the TOM as a measure of similarity from 200 simulation runs. Vertical panels represent varying correlation between active clusters. Horizontal panels represent different signal-to-noise ratios.




Figure S11: Simulation 1 – Correct Sparsity based on the training set using the TOM as a measure of similarity from 200 simulation runs. Vertical panels represent varying correlation between active clusters. Horizontal panels represent different signal-to-noise ratios.

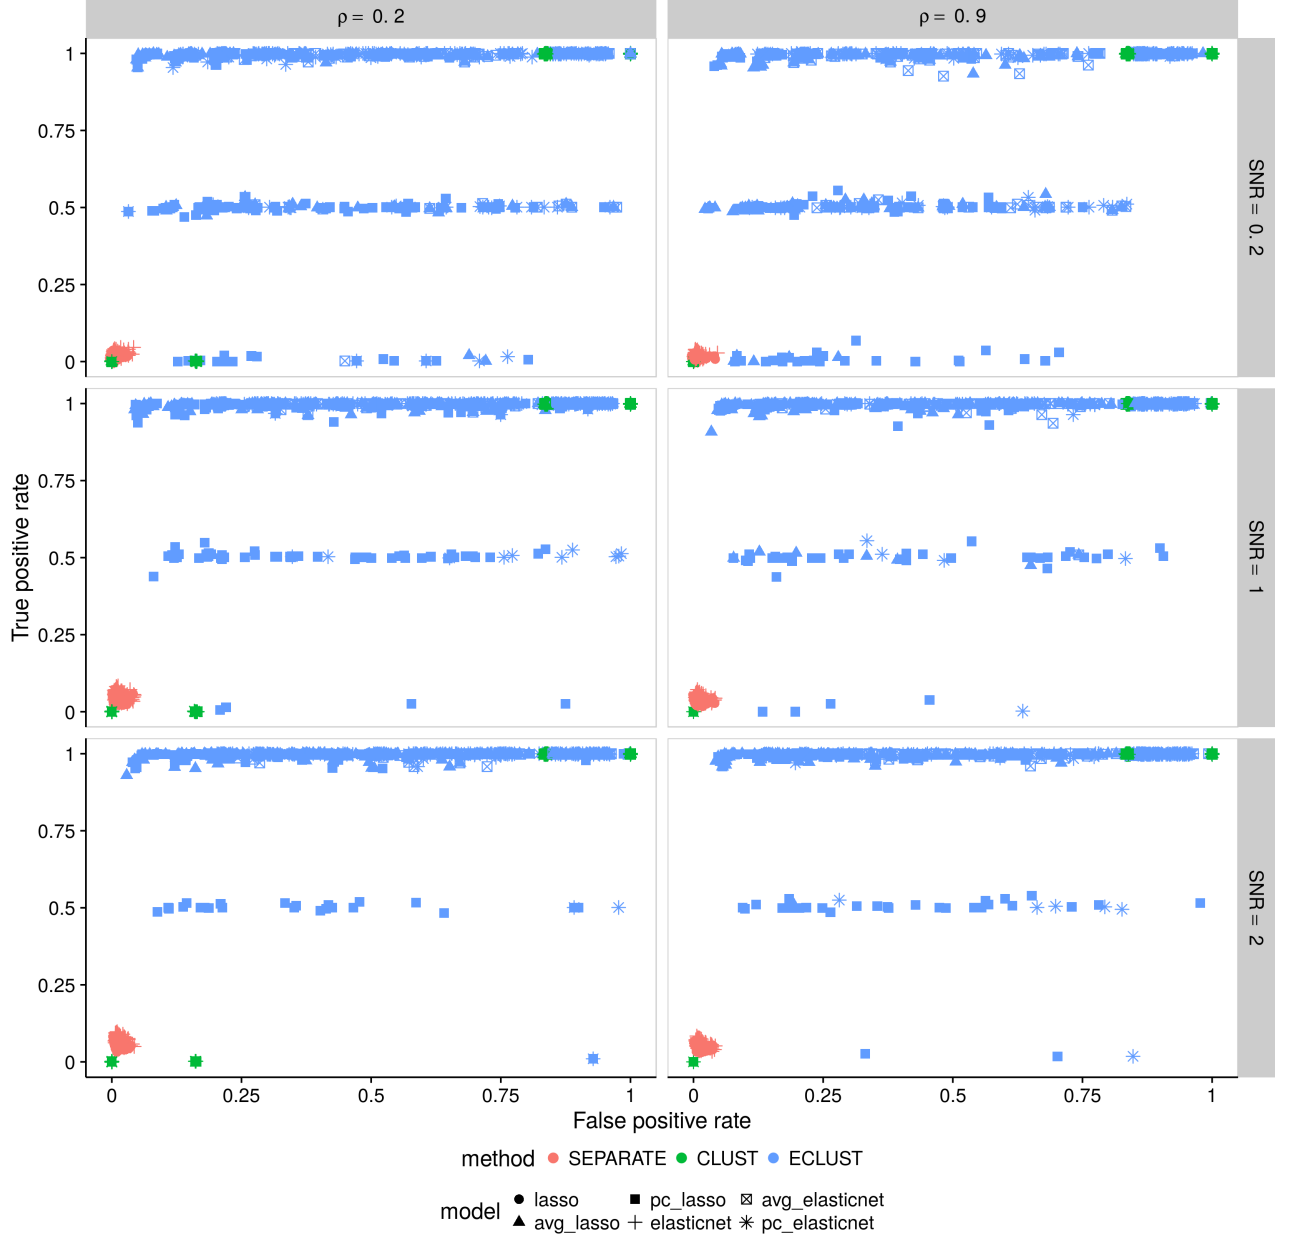


Figure S12: Simulation 1 – True positive rate vs. false positive rate based on the training set using the TOM as a measure of similarity. Each point represents 1 simulation run (there are a total of 200 simulation runs). Vertical panels represent varying correlation between active clusters. Horizontal panels represent different signal-to-noise ratios.




Figure S13: Simulation 1 – Average Jaccard Index from 10 CV folds of the training set using the TOM as a measure of similarity. We fit the model to each of the 10 CV folds resulting in 10 sets of selected predictors. We then calculate the Jaccard Index between all $\binom{10}{2}$ possible combinations of these sets and take the average. This process is repeated for each of the 200 simulation runs. Vertical panels represent varying correlation between active clusters. Horizontal panels represent different signal-to-noise ratios.




Figure S14: Simulation 1 – Average Spearman correlation from 10 CV folds of the training set using the TOM as a measure of similarity. We fit the model to each of the 10 CV folds resulting in 10 sets of estimated regression coefficients. We then calculate the Spearman correlation between all $\binom{10}{2}$ possible combinations of these sets and take the average. This process is repeated for each of the 200 simulation runs. Vertical panels represent varying correlation between active clusters. Horizontal panels represent different signal-to-noise ratios.




Figure S15: Simulation 1 – Average Pearson correlation from 10 CV folds of the training set using the TOM as a measure of similarity. We fit the model to each of the 10 CV folds resulting in 10 sets of estimated regression coefficients. We then calculate the Pearson correlation between all $\binom{10}{2}$ possible combinations of these sets and take the average. This process is repeated for each of the 200 simulation runs. Vertical panels represent varying correlation between active clusters. Horizontal panels represent different signal-to-noise ratios.

D.2 Simulation 2

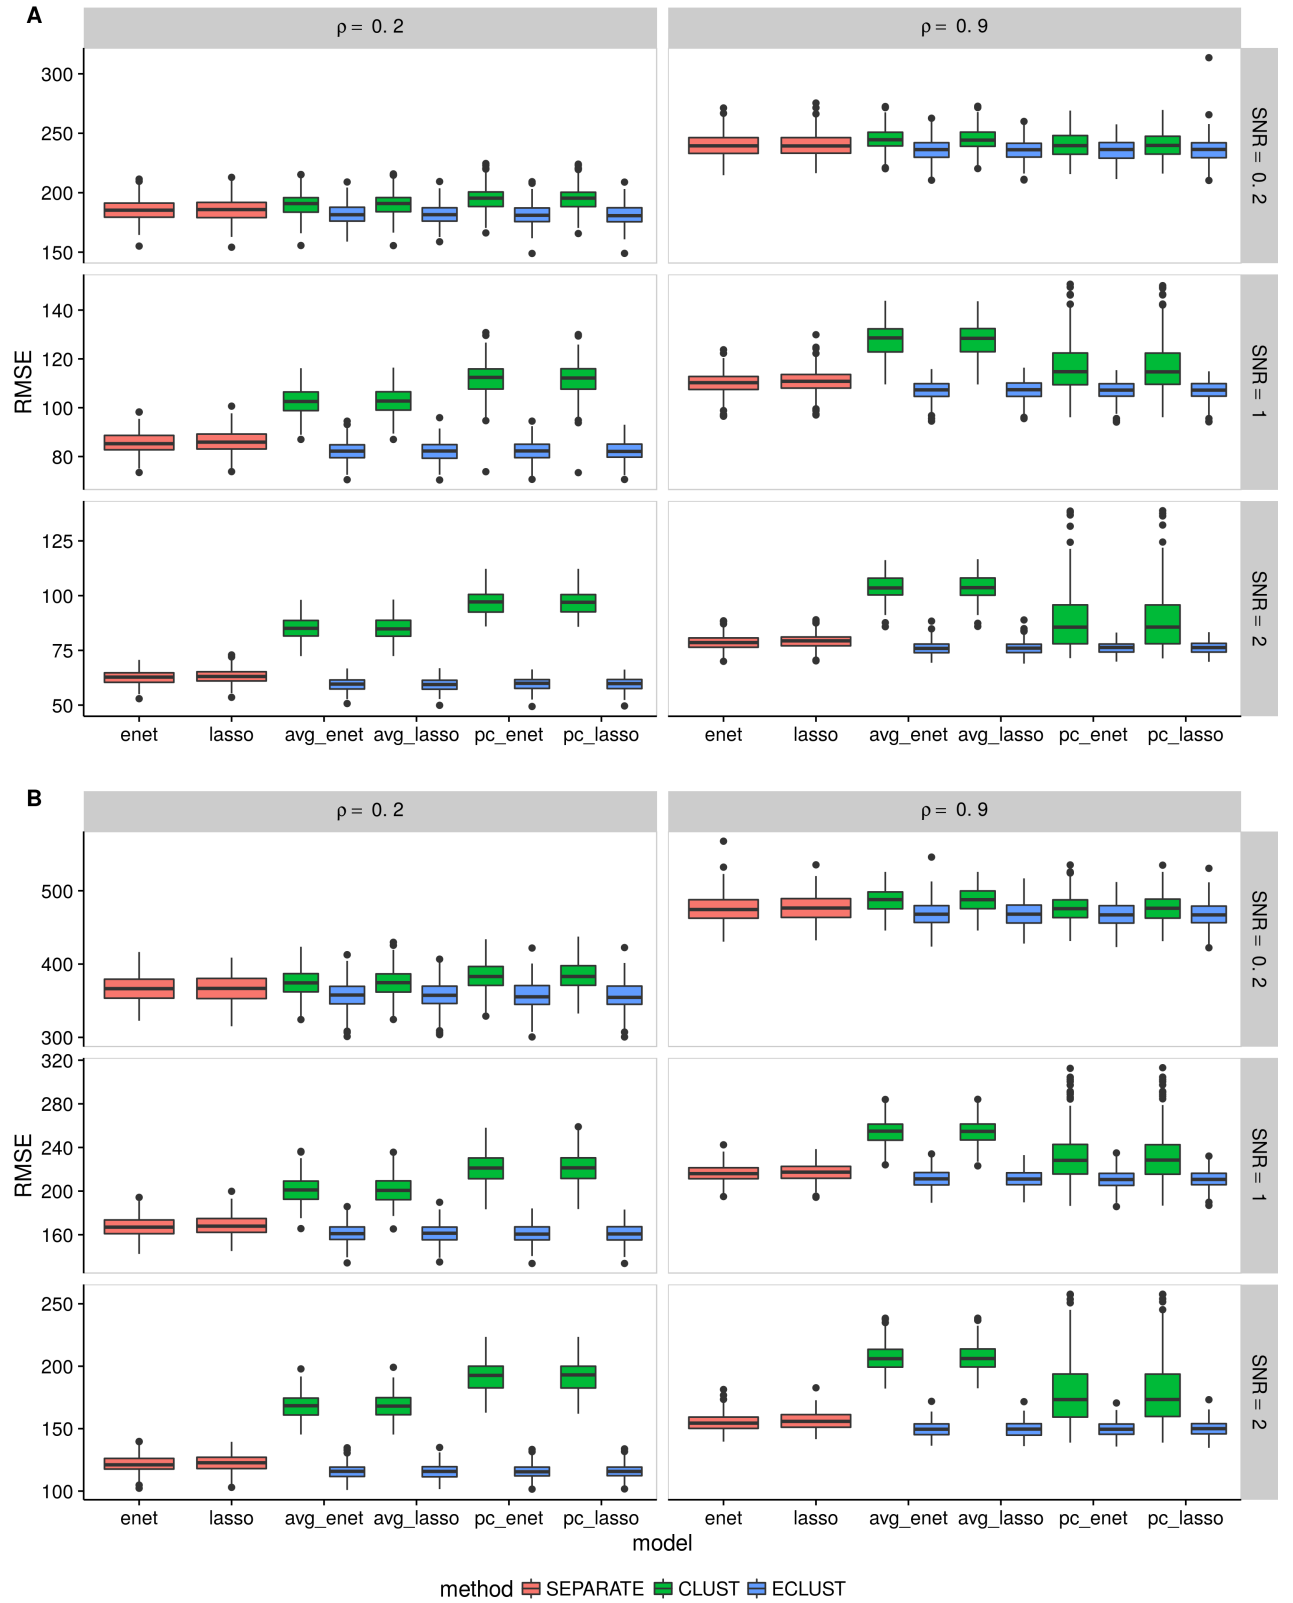


Figure S16: Simulation 2 – Root mean squared error on an independent test set using the TOM as a measure of similarity from 200 simulation runs. (A) $\alpha_j \sim \text{Unif}[0.4, 0.6]$, (B) $\alpha_j \sim \text{Unif}[1.9, 2.1]$. Vertical panels represent varying correlation between active clusters. Horizontal panels represent different signal-to-noise ratios.

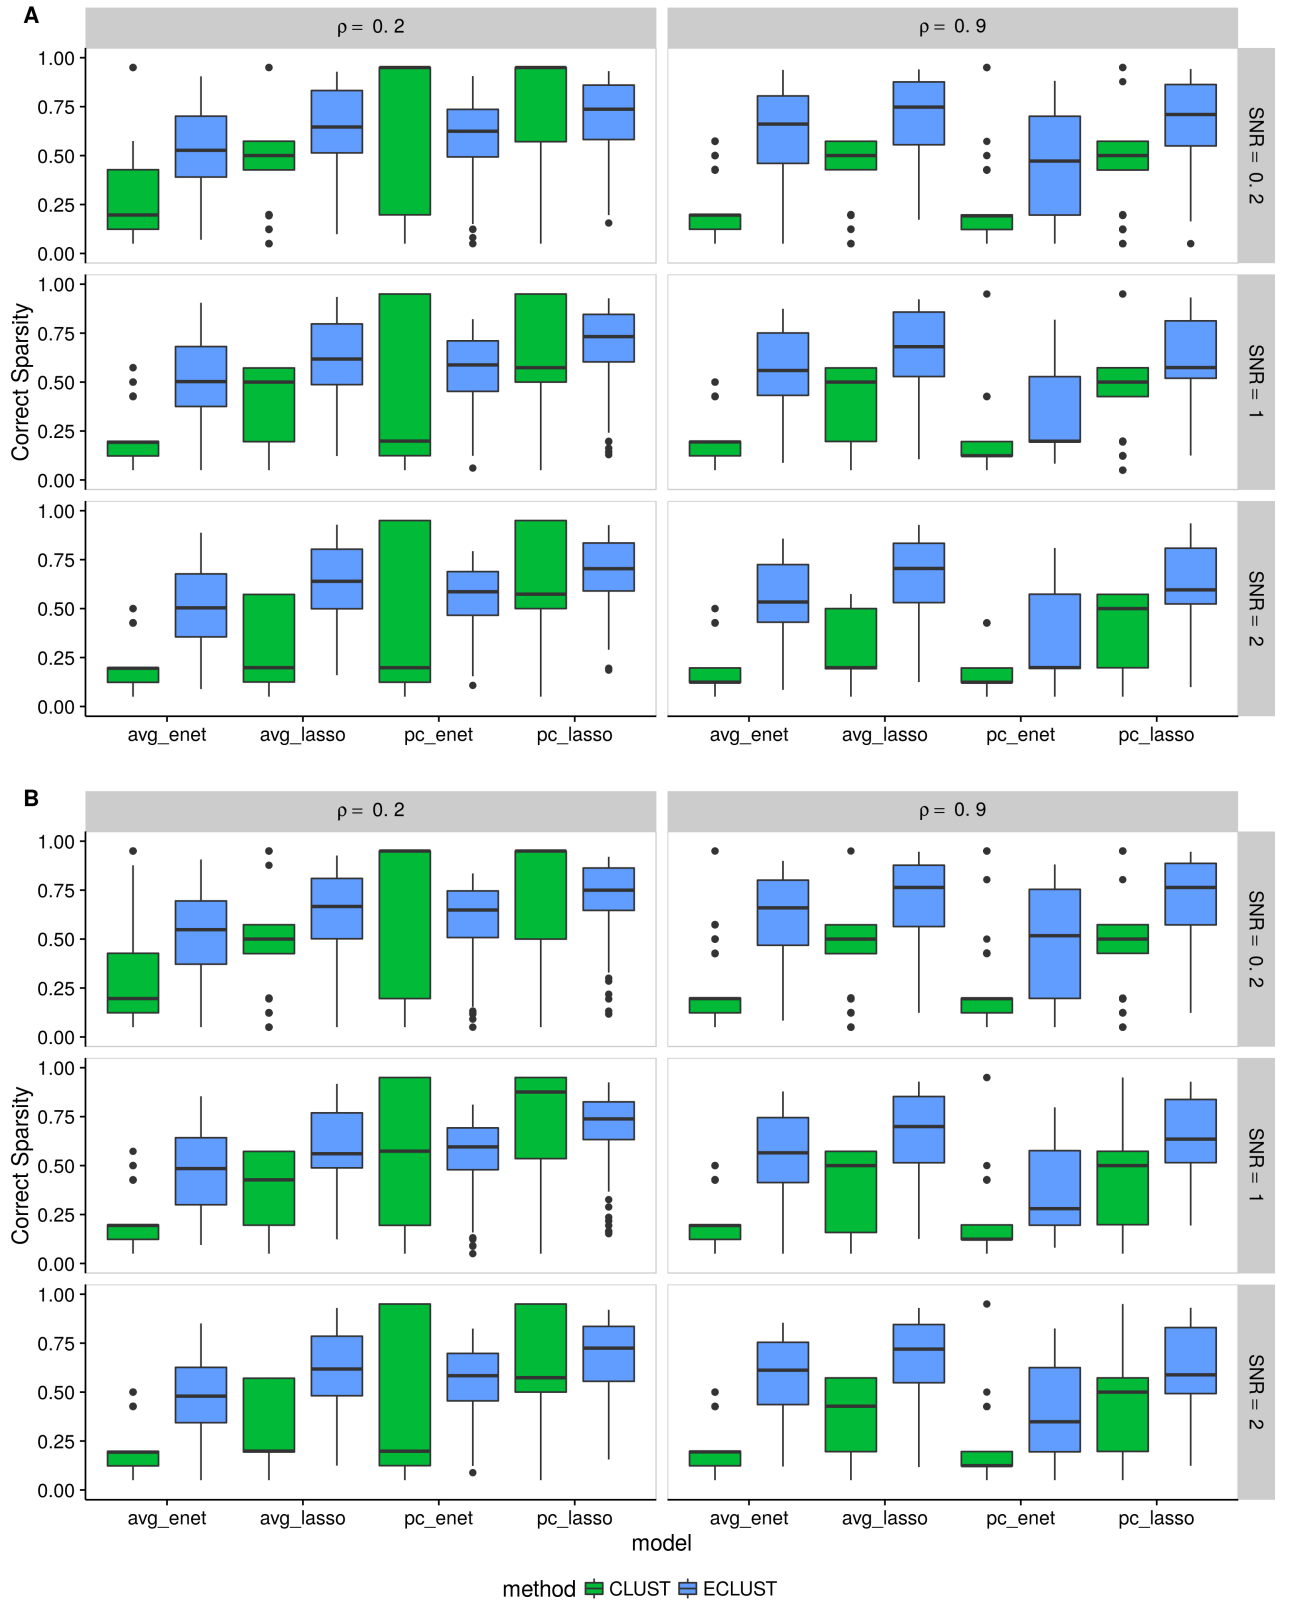


Figure S17: Simulation 2 – Correct Sparsity based on the training set using the TOM as a measure of similarity from 200 simulation runs. (A) $\alpha_j \sim \text{Unif}[0.4, 0.6]$, (B) $\alpha_j \sim \text{Unif}[1.9, 2.1]$. Vertical panels represent varying correlation between active clusters. Horizontal panels represent different signal-to-noise ratios.

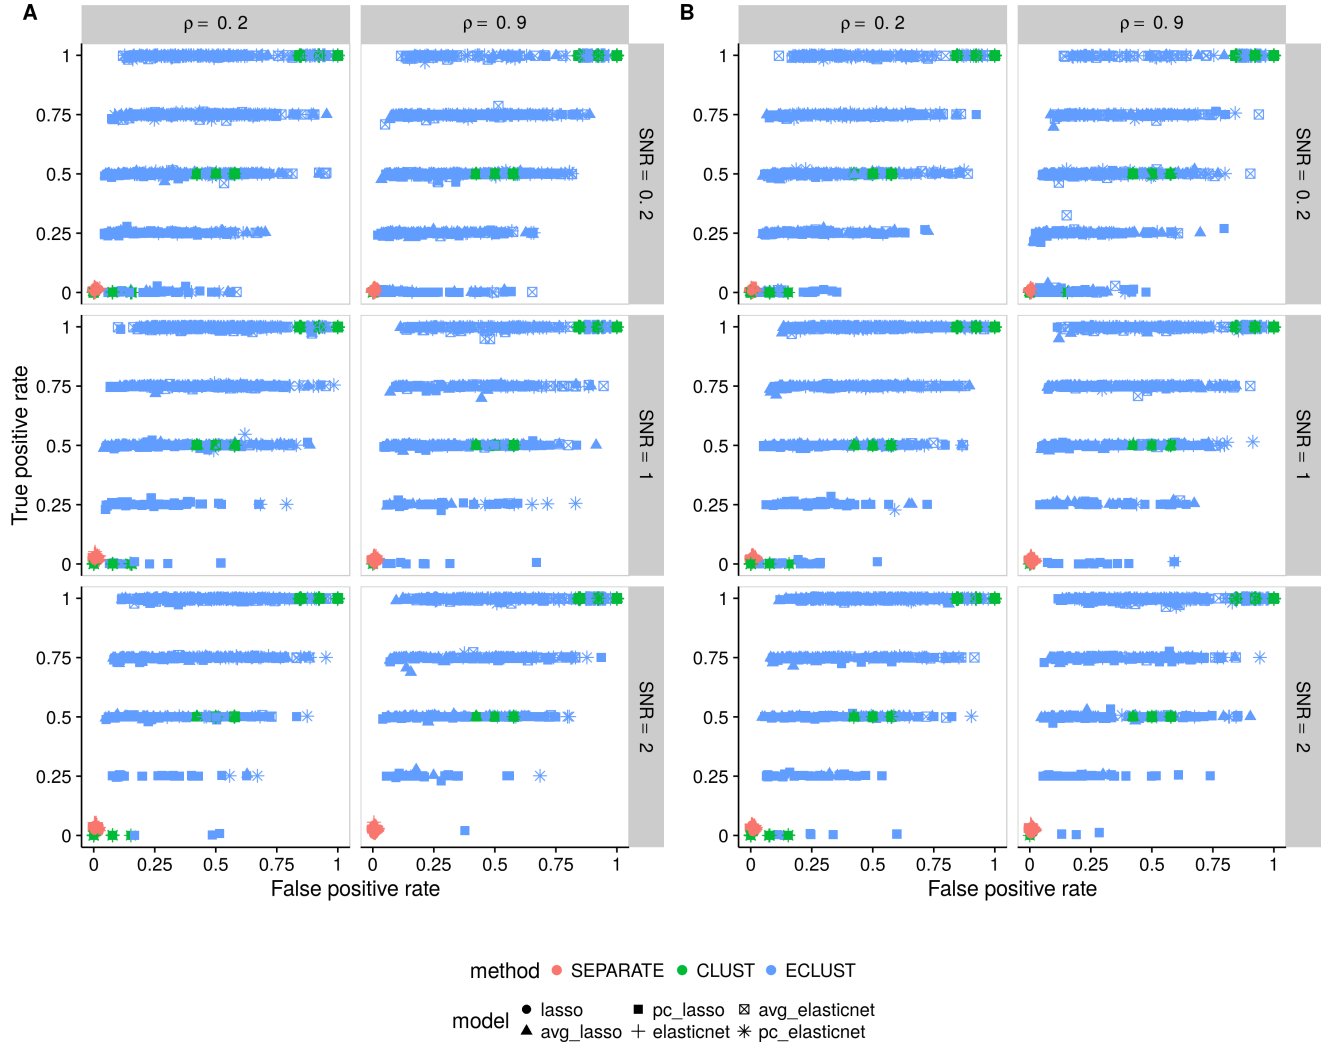


Figure S18: Simulation 2 – True positive rate vs. false positive rate based on the training set using the TOM as a measure of similarity. (A) $\alpha_j \sim \text{Unif}[0.4, 0.6]$, (B) $\alpha_j \sim \text{Unif}[1.9, 2.1]$. Each point represents 1 simulation run (there are a total of 200 simulation runs). Vertical panels represent varying correlation between active clusters. Horizontal panels represent different signal-to-noise ratios.

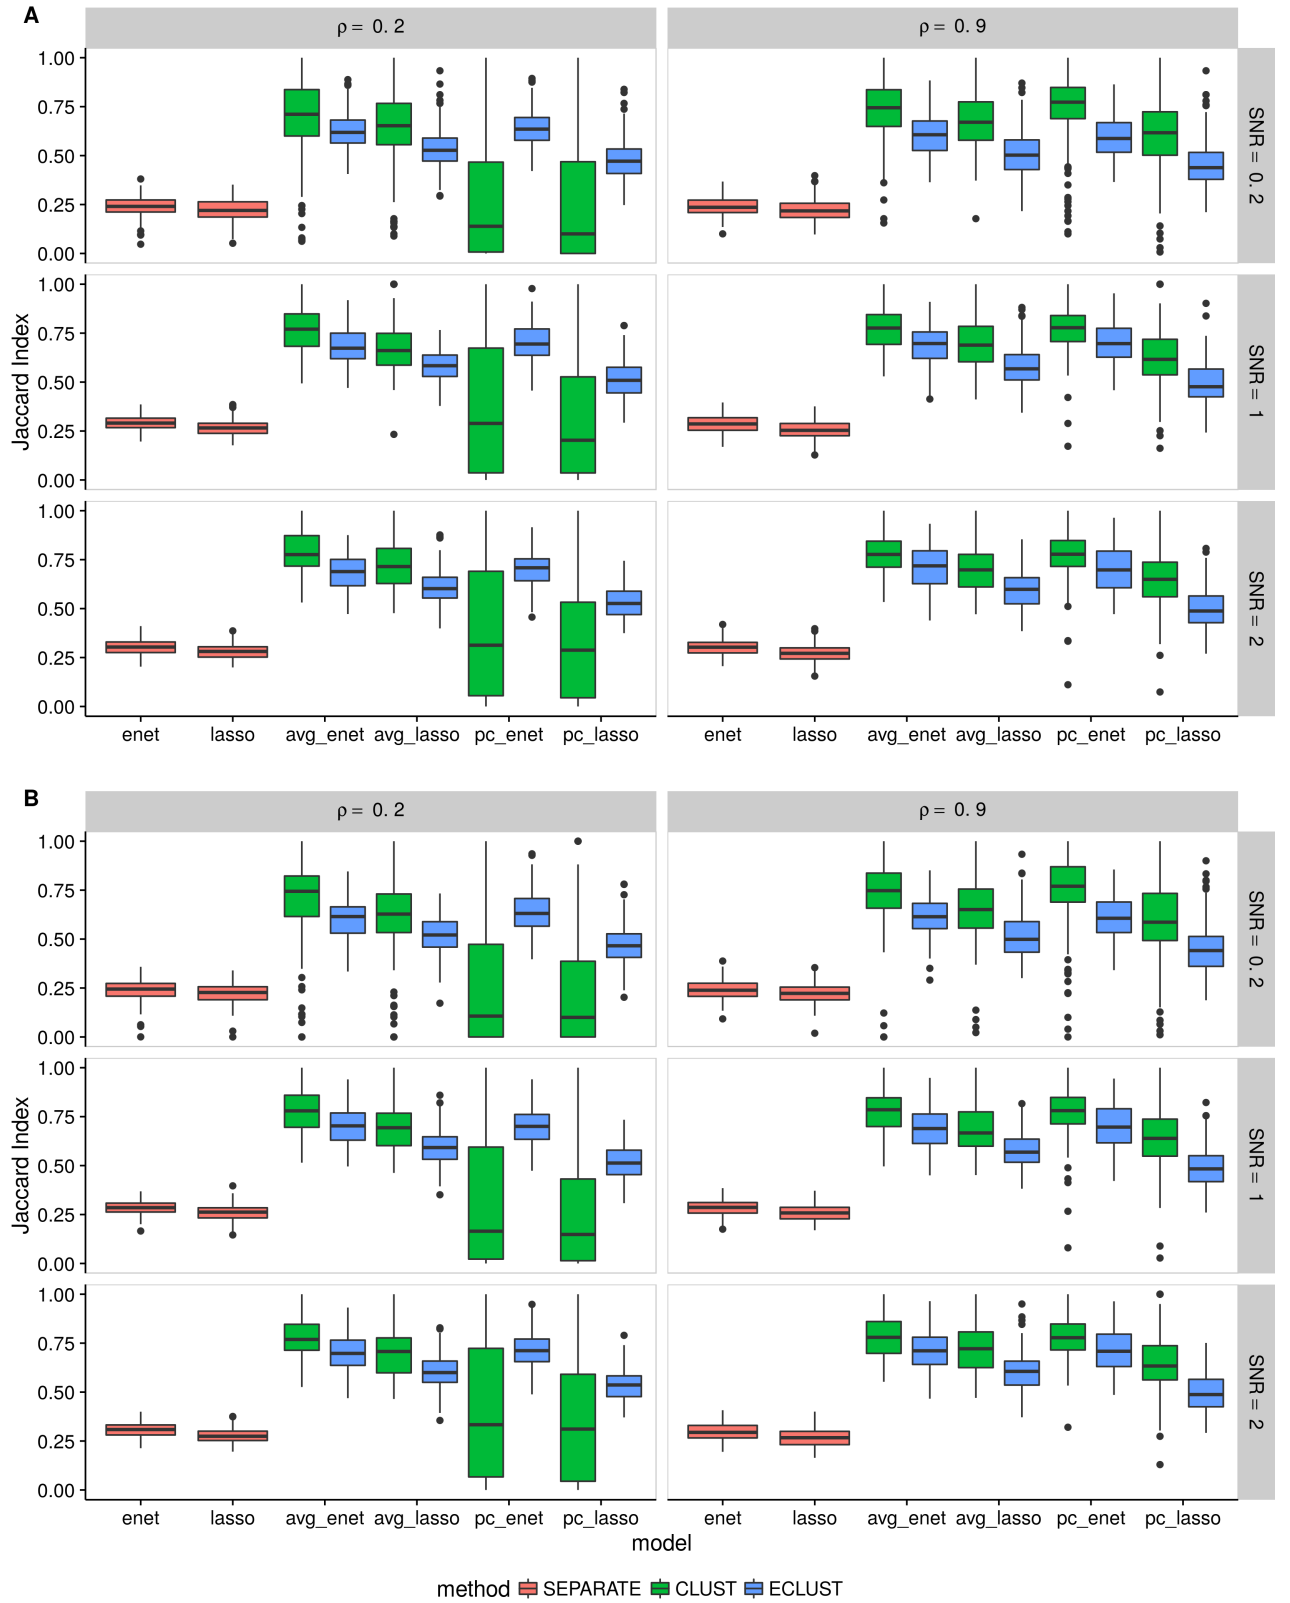


Figure S19: Simulation 2 – Average Jaccard Index from 10 CV folds of the training set using the TOM as a measure of similarity. (A) $\alpha_j \sim \text{Unif}[0.4, 0.6]$, (B) $\alpha_j \sim \text{Unif}[1.9, 2.1]$. We fit the model to each of the 10 CV folds resulting in 10 sets of selected predictors. We then calculate the Jaccard Index between all $\binom{10}{2}$ possible combinations of these sets and take the average. This process is repeated for each of the 200 simulation runs. Vertical panels represent varying correlation between active clusters. Horizontal panels represent different signal-to-noise ratios.

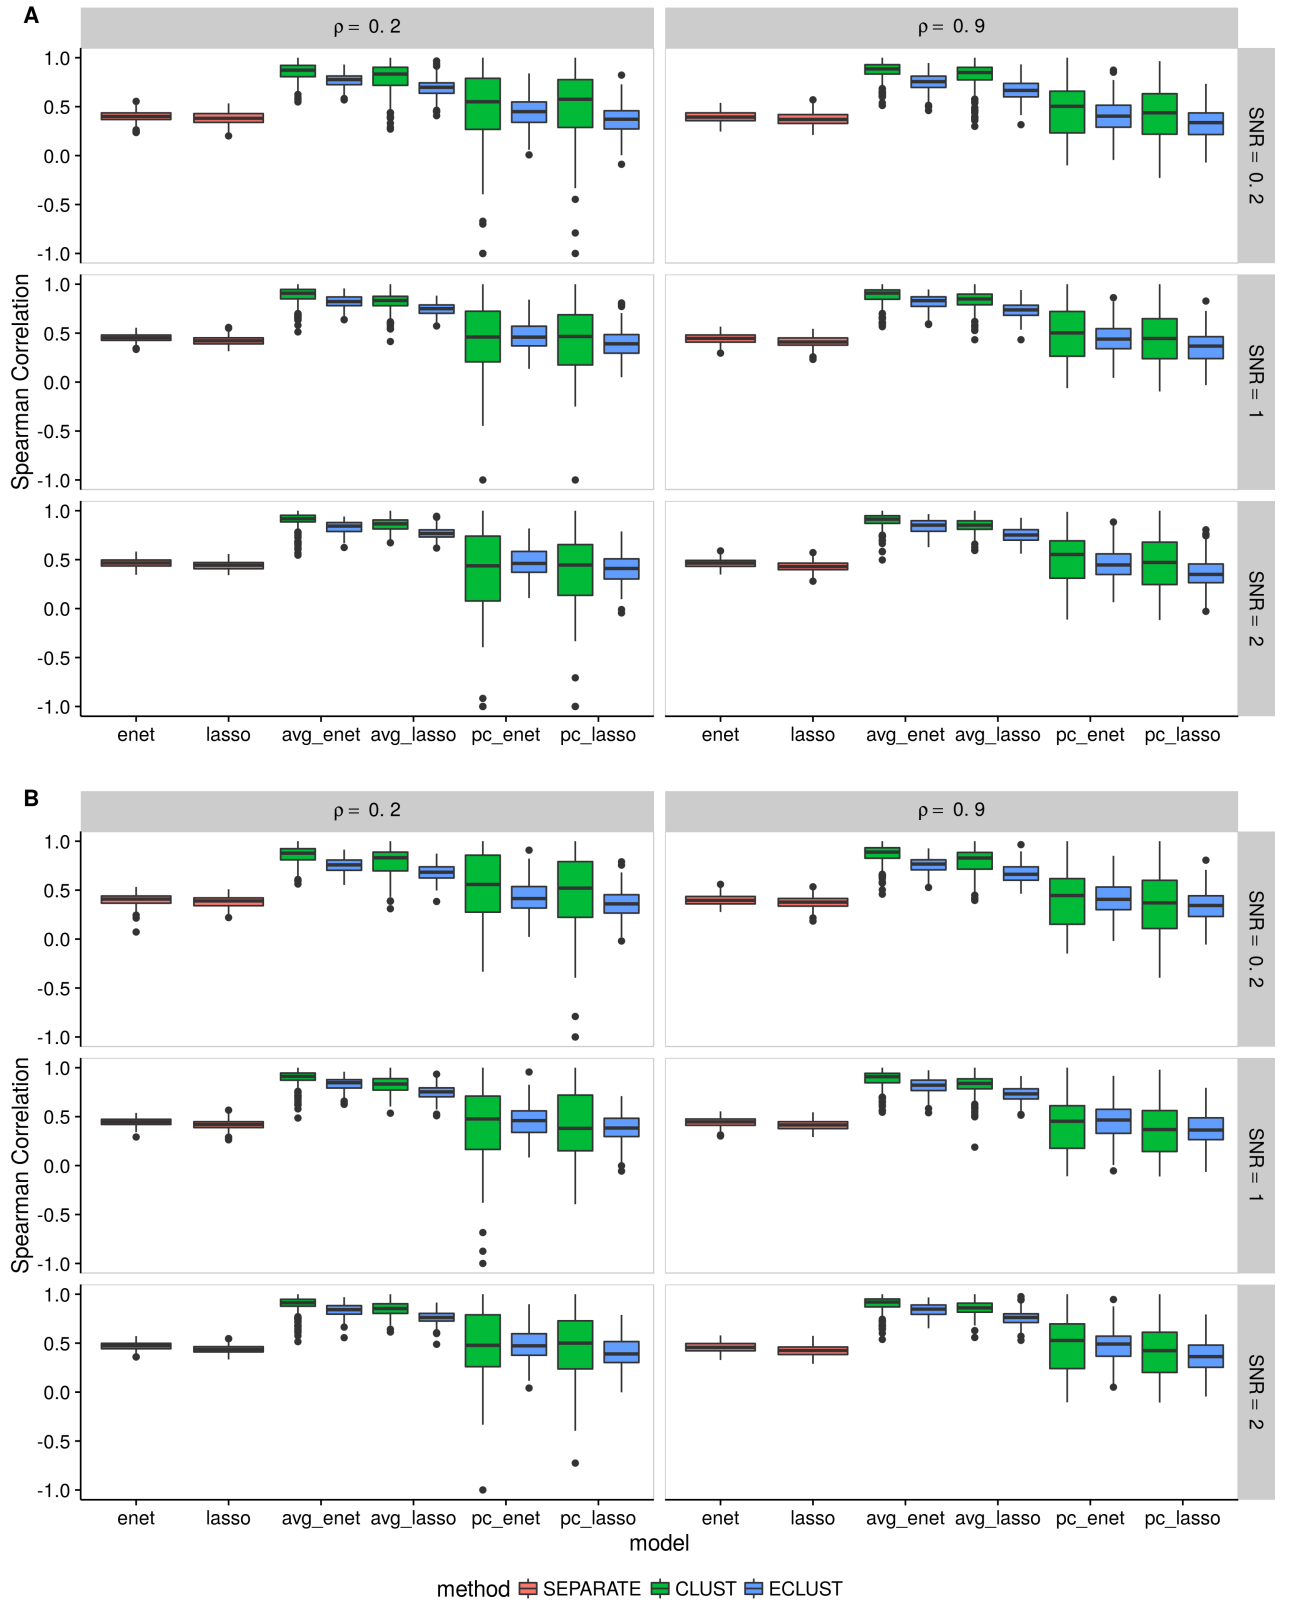


Figure S20: Simulation 2 – Average Spearman correlation from 10 CV folds of the training set using the TOM as a measure of similarity. (A) $\alpha_j \sim \text{Unif}[0.4, 0.6]$, (B) $\alpha_j \sim \text{Unif}[1.9, 2.1]$. We fit the model to each of the 10 CV folds resulting in 10 sets of estimated regression coefficients. We then calculate the Spearman correlation between all $\binom{10}{2}$ possible combinations of these sets and take the average. This process is repeated for each of the 200 simulation runs. Vertical panels represent varying correlation between active clusters. Horizontal panels represent different signal-to-noise ratios.




Figure S21: Simulation 2 – Average Pearson correlation from 10 CV folds of the training set using the TOM as a measure of similarity. (A) $\alpha_j \sim \text{Unif}[0.4, 0.6]$, (B) $\alpha_j \sim \text{Unif}[1.9, 2.1]$. We fit the model to each of the 10 CV folds resulting in 10 sets of estimated regression coefficients. We then calculate the Pearson correlation between all $\binom{10}{2}$ possible combinations of these sets and take the average. This process is repeated for each of the 200 simulation runs. Vertical panels represent varying correlation between active clusters. Horizontal panels represent different signal-to-noise ratios.

D.3 Simulation 3




Figure S22: Simulation 3 – Root mean squared error on an independent test set using the TOM as a measure of similarity from 200 simulation runs. Vertical panels represent varying correlation between active clusters. Horizontal panels represent different signal-to-noise ratios.




Figure S23: Simulation 3 – Correct Sparsity based on the training set using the TOM as a measure of similarity from 200 simulation runs. Vertical panels represent varying correlation between active clusters. Horizontal panels represent different signal-to-noise ratios.

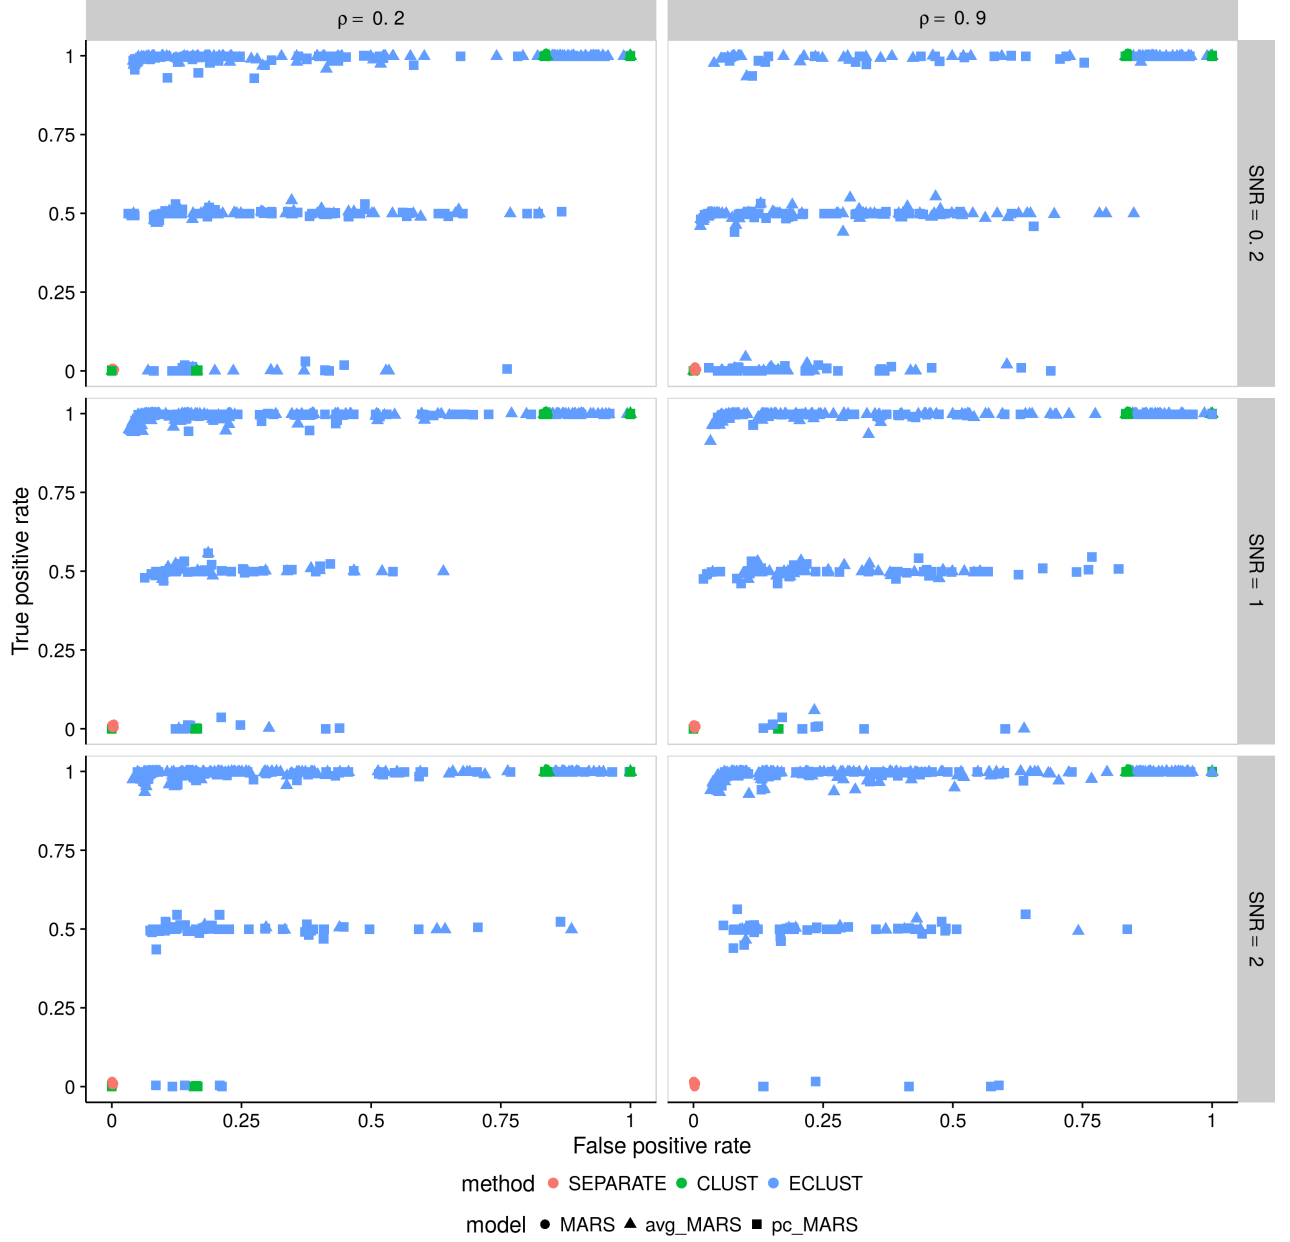


Figure S24: Simulation 3 – True positive rate vs. false positive rate based on the training set using the TOM as a measure of similarity. Each point represents 1 simulation run (there are a total of 200 simulation runs). Vertical panels represent varying correlation between active clusters. Horizontal panels represent different signal-to-noise ratios.




Figure S25: Simulation 3 – Average Jaccard Index from 10 CV folds of the training set using the TOM as a measure of similarity. We fit the model to each of the 10 CV folds resulting in 10 sets of selected predictors. We then calculate the Jaccard Index between all $\binom{10}{2}$ possible combinations of these sets and take the average. This process is repeated for each of the 200 simulation runs. Vertical panels represent varying correlation between active clusters. Horizontal panels represent different signal-to-noise ratios.

E Simulation Results Using Pearson Correlations as a Measure of Similarity

E.1 Simulation 1

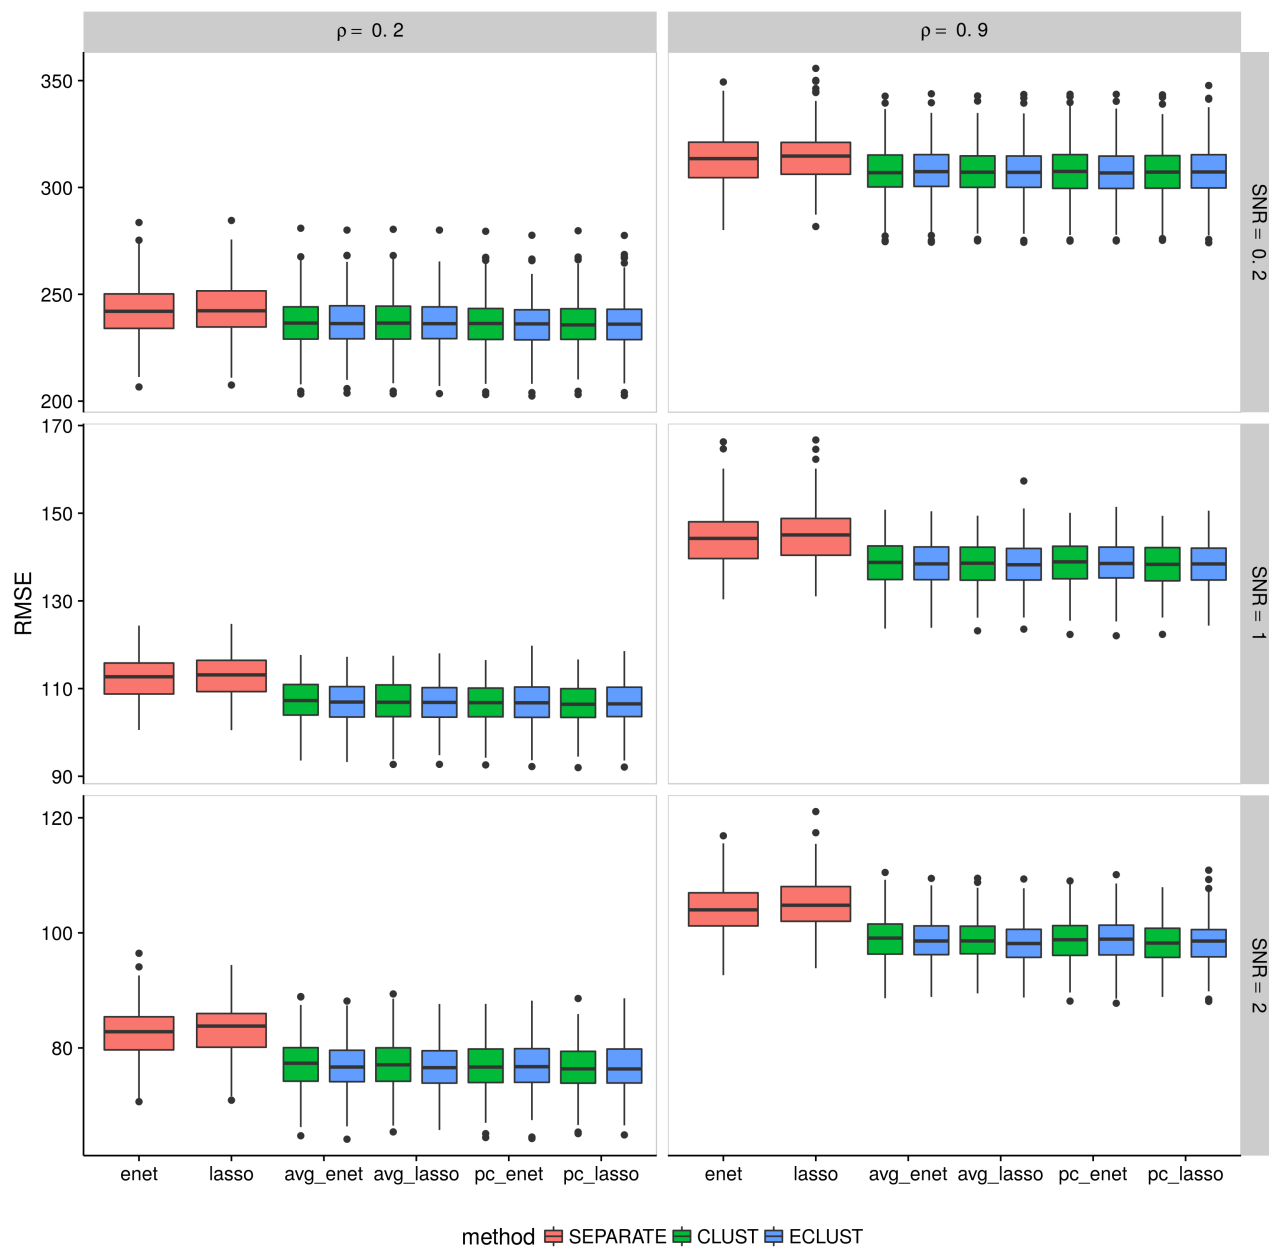


Figure S26: Simulation 1 – Root mean squared error on an independent test set using the Correlation as a measure of similarity from 200 simulation runs. Vertical panels represent varying correlation between active clusters. Horizontal panels represent different signal-to-noise ratios.




Figure S27: Simulation 1 – Correct Sparsity based on the training set using the Pearson correlation as a measure of similarity from 200 simulation runs. Vertical panels represent varying correlation between active clusters. Horizontal panels represent different signal-to-noise ratios.




Figure S28: Simulation 1 – True positive rate vs. false positive rate based on the training set using the Pearson correlation as a measure of similarity. Each point represents 1 simulation run (there are a total of 200 simulation runs). Vertical panels represent varying correlation between active clusters. Horizontal panels represent different signal-to-noise ratios.




Figure S29: Simulation 1 – Average Jaccard Index from 10 CV folds of the training set using the Pearson correlation as a measure of similarity. We fit the model to each of the 10 CV folds resulting in 10 sets of selected predictors. We then calculate the Jaccard Index between all $\binom{10}{2}$ possible combinations of these sets and take the average. This process is repeated for each of the 200 simulation runs. Vertical panels represent varying correlation between active clusters. Horizontal panels represent different signal-to-noise ratios.




Figure S30: Simulation 1 – Average Spearman correlation from 10 CV folds of the training set using the Pearson correlation as a measure of similarity. We fit the model to each of the 10 CV folds resulting in 10 sets of estimated regression coefficients. We then calculate the Spearman correlation between all $\binom{10}{2}$ possible combinations of these sets and take the average. This process is repeated for each of the 200 simulation runs. Vertical panels represent varying correlation between active clusters. Horizontal panels represent different signal-to-noise ratios.




Figure S31: Simulation 1 – Average Pearson correlation from 10 CV folds of the training set using the Pearson correlation as a measure of similarity. We fit the model to each of the 10 CV folds resulting in 10 sets of estimated regression coefficients. We then calculate the Pearson correlation between all $\binom{10}{2}$ possible combinations of these sets and take the average. This process is repeated for each of the 200 simulation runs. Vertical panels represent varying correlation between active clusters. Horizontal panels represent different signal-to-noise ratios.

E.2 Simulation 2

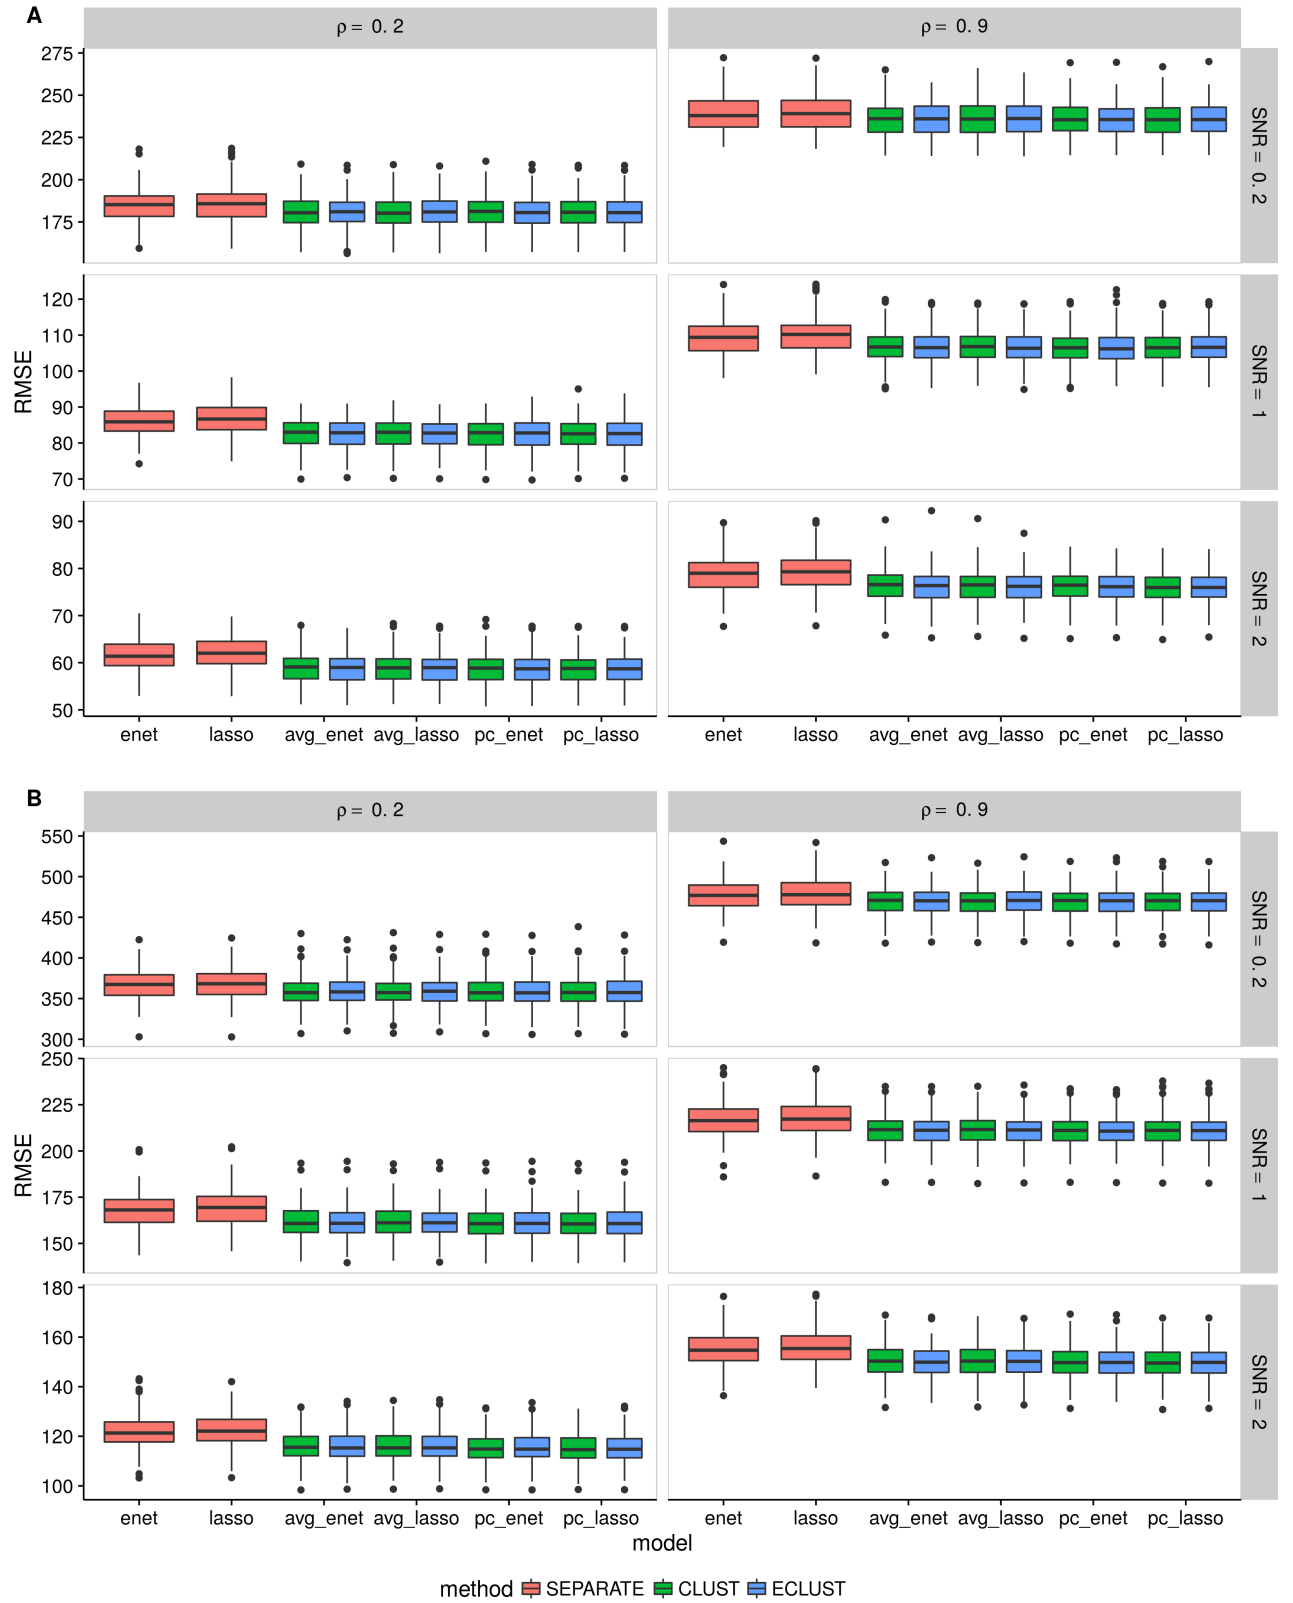


Figure S32: Simulation 2 – Root mean squared error on an independent test set using the Pearson correlation as a measure of similarity from 200 simulation runs. (A) $\alpha_j \sim \text{Unif}[0.4, 0.6]$, (B) $\alpha_j \sim \text{Unif}[1.9, 2.1]$. Vertical panels represent varying correlation between active clusters. Horizontal panels represent different signal-to-noise ratios.

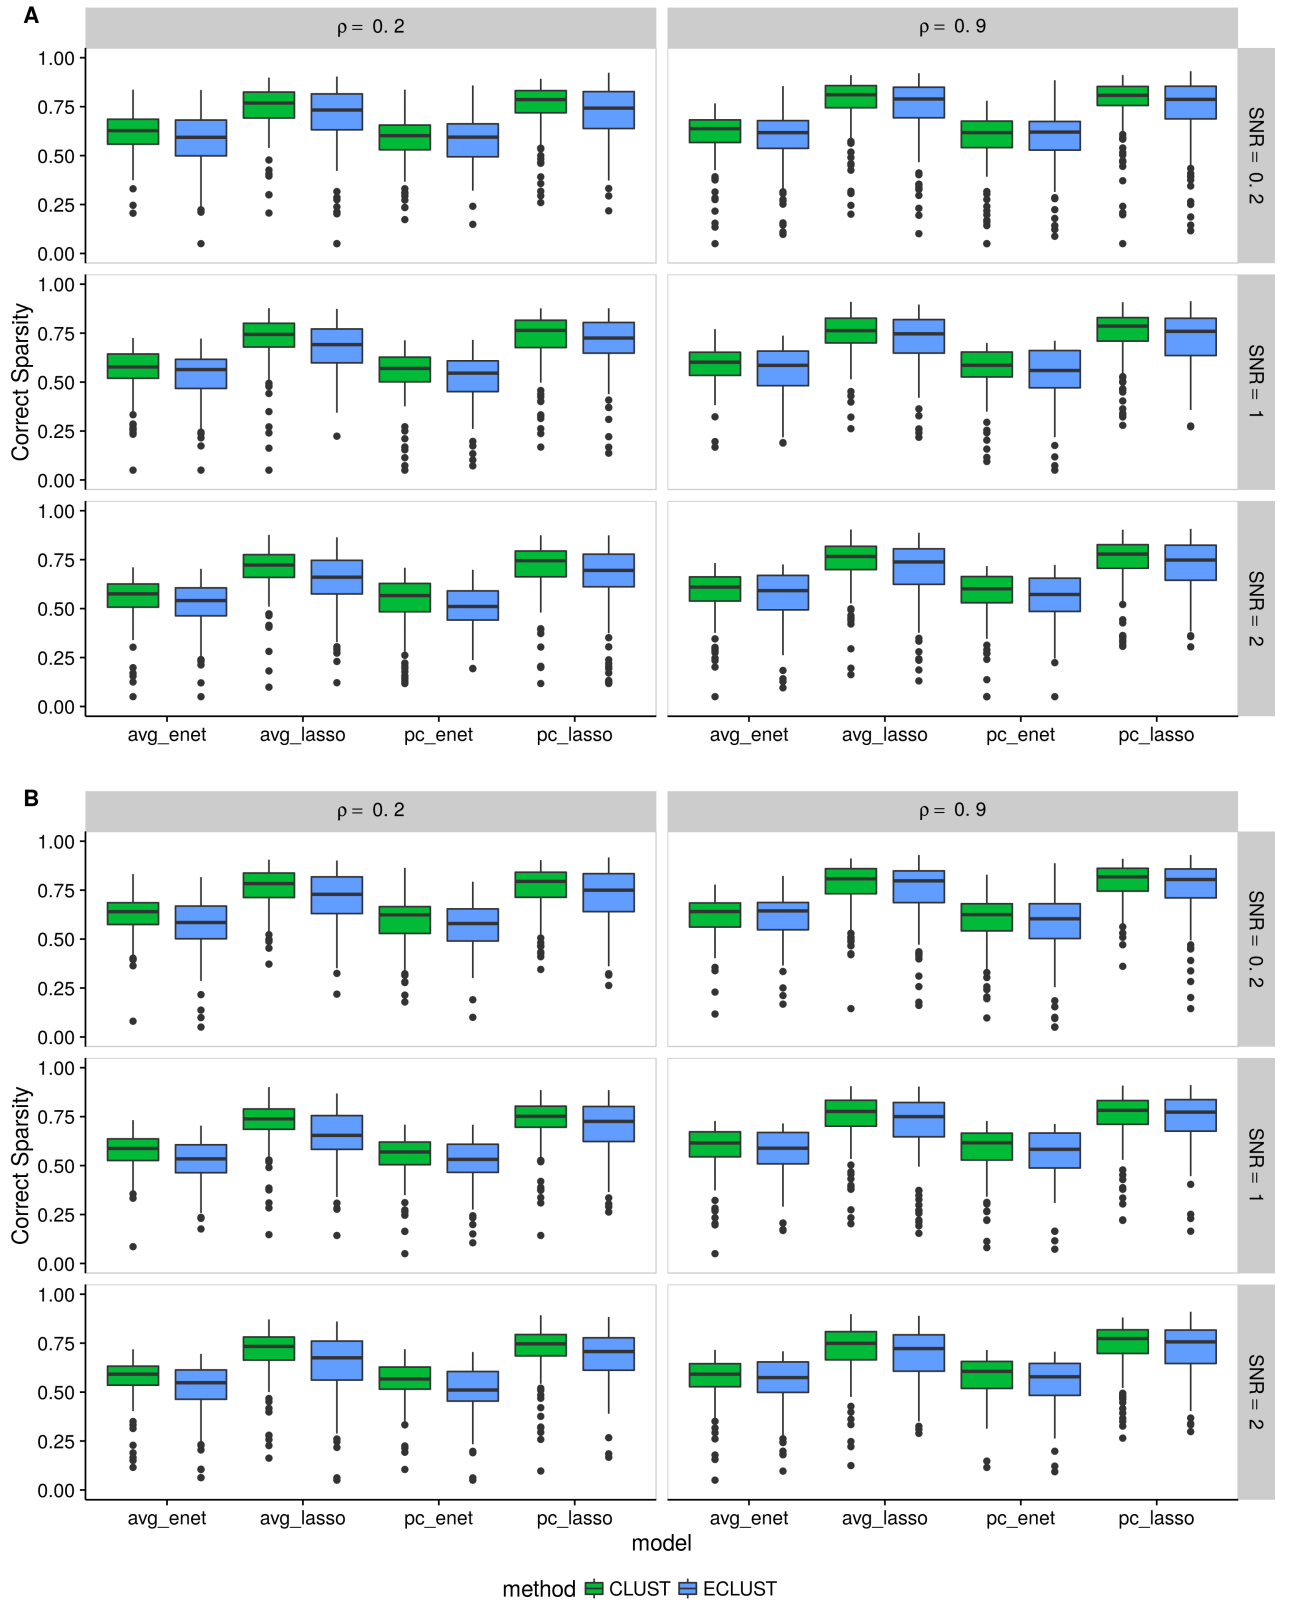


Figure S33: Simulation 2 – Correct Sparsity based on the training set using the Pearson correlation as a measure of similarity from 200 simulation runs. (A) $\alpha_j \sim \text{Unif}[0.4, 0.6]$, (B) $\alpha_j \sim \text{Unif}[1.9, 2.1]$. Vertical panels represent varying correlation between active clusters. Horizontal panels represent different signal-to-noise ratios.

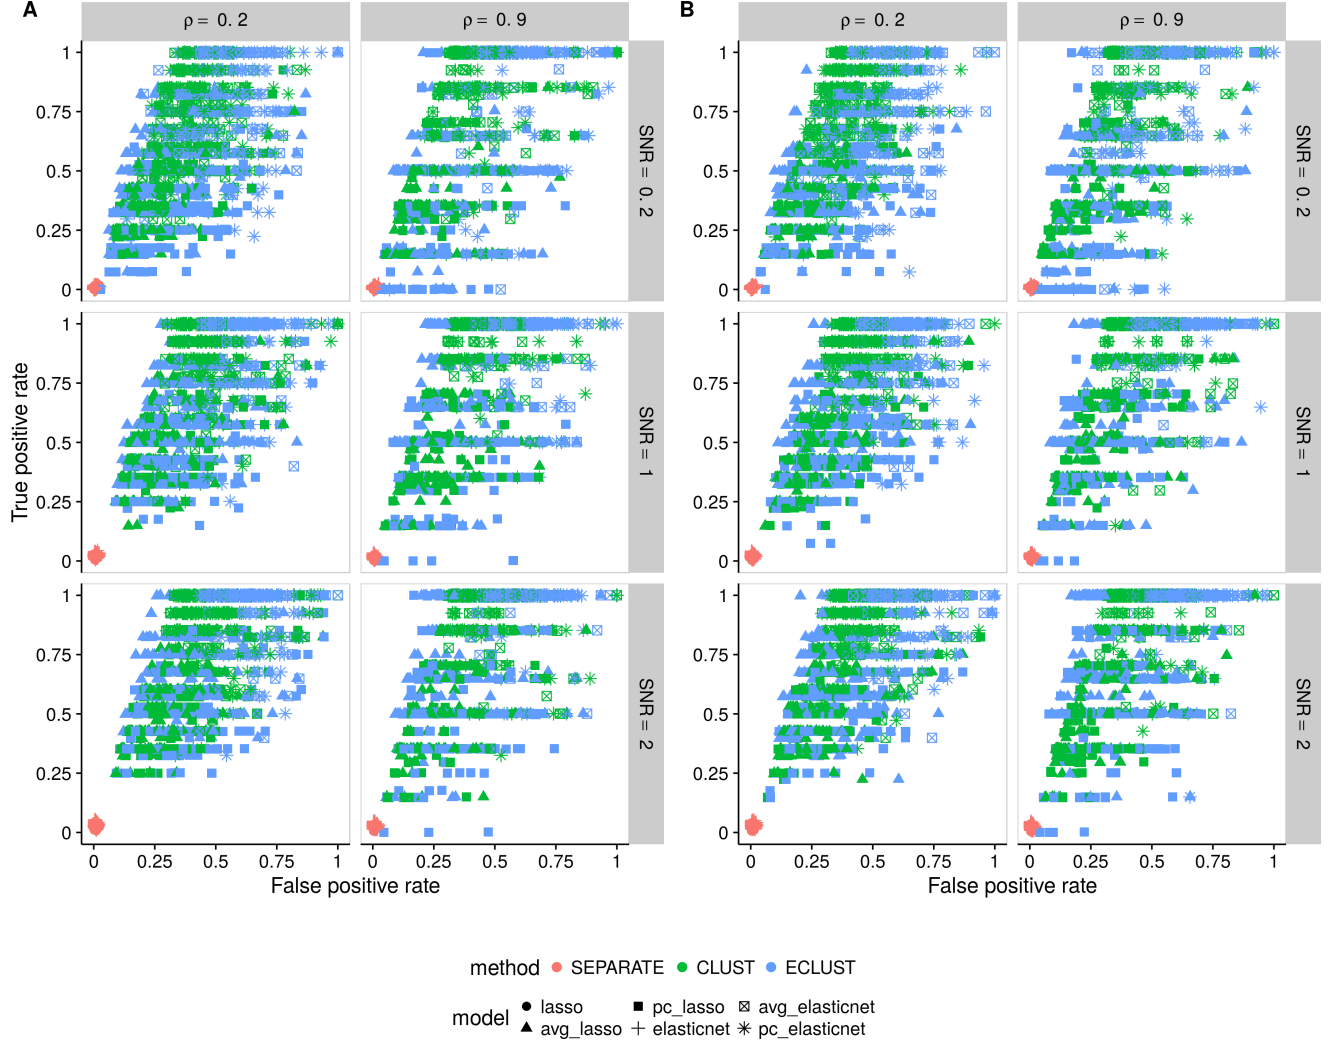


Figure S34: Simulation 2 – True positive rate vs. false positive rate based on the training set using the Pearson correlation as a measure of similarity. (A) $\alpha_j \sim \text{Unif}[0.4, 0.6]$, (B) $\alpha_j \sim \text{Unif}[1.9, 2.1]$. Each point represents 1 simulation run (there are a total of 200 simulation runs). Vertical panels represent varying correlation between active clusters. Horizontal panels represent different signal-to-noise ratios.

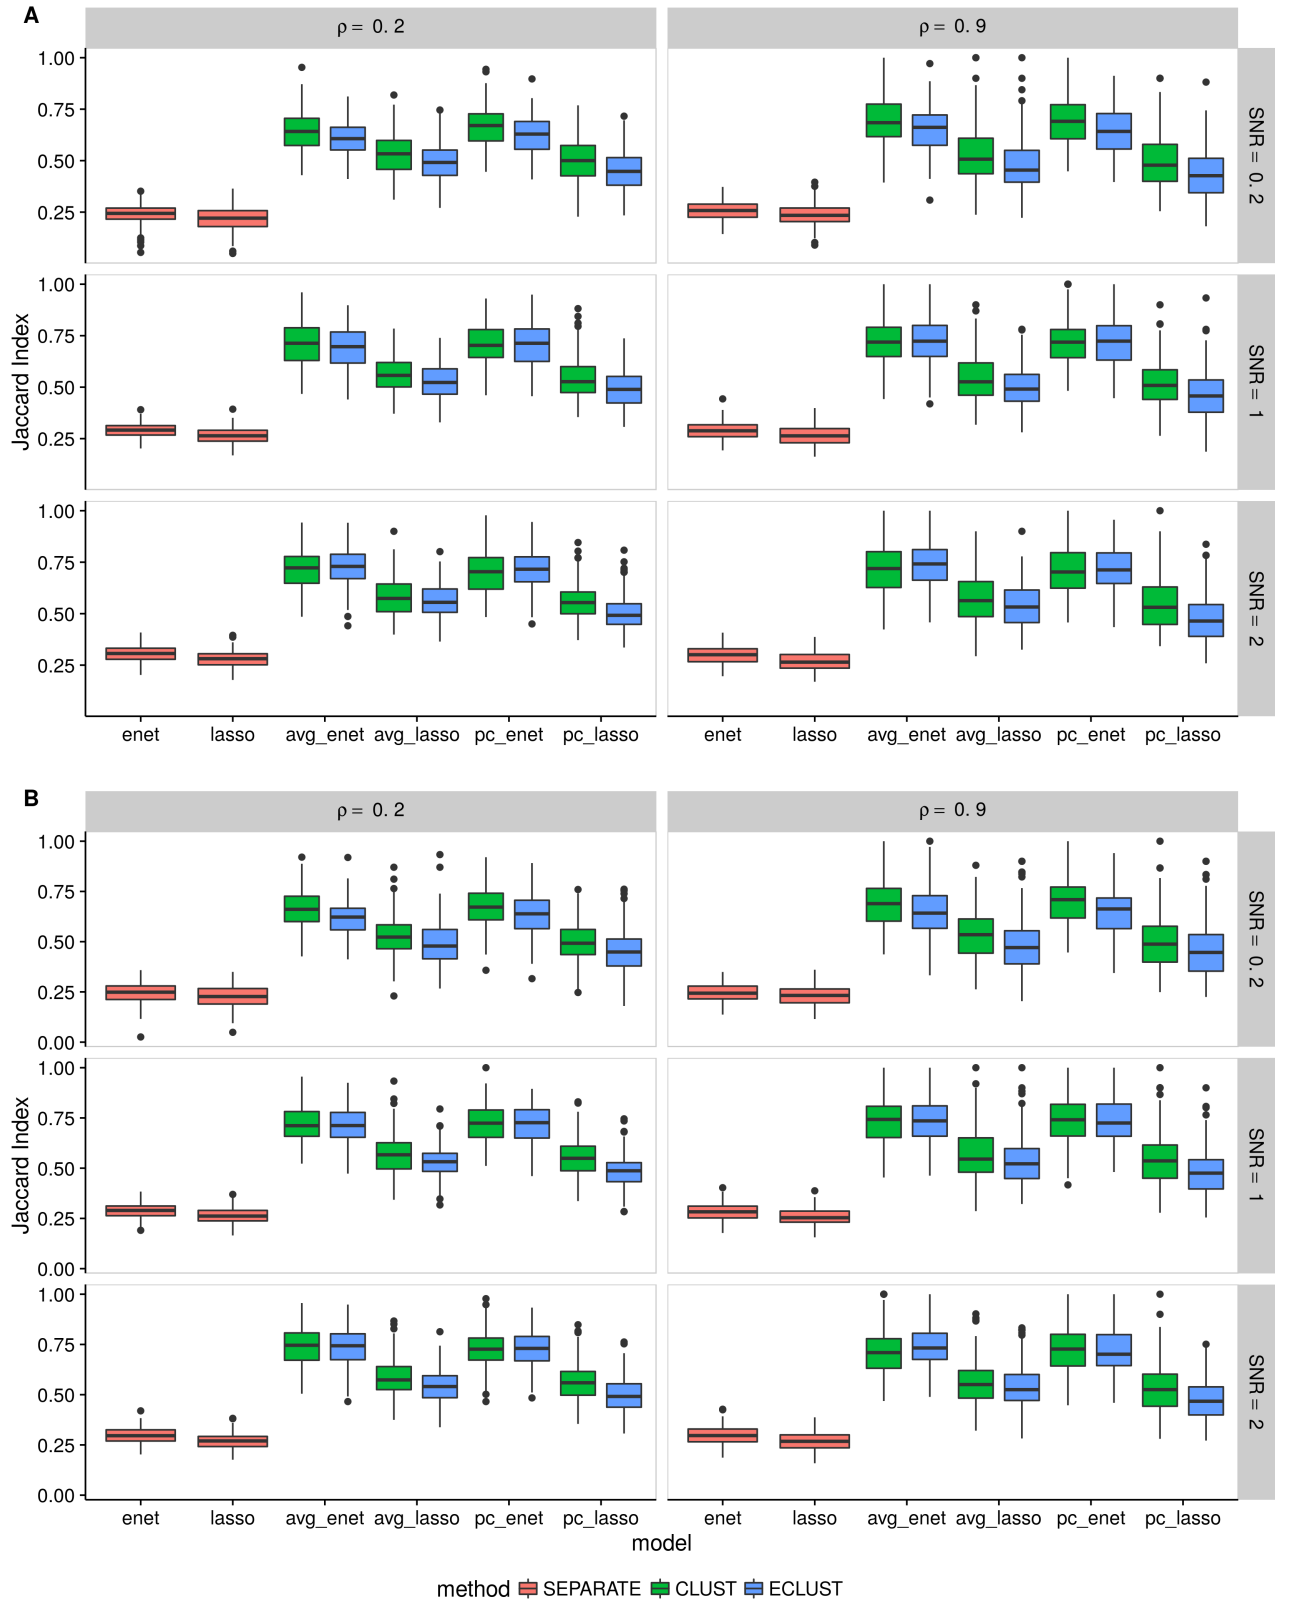


Figure S35: Simulation 2 – Average Jaccard Index from 10 CV folds of the training set using the Pearson correlation as a measure of similarity. (A) $\alpha_j \sim \text{Unif}[0.4, 0.6]$, (B) $\alpha_j \sim \text{Unif}[1.9, 2.1]$. We fit the model to each of the 10 CV folds resulting in 10 sets of selected predictors. We then calculate the Jaccard Index between all $\binom{10}{2}$ possible combinations of these sets and take the average. This process is repeated for each of the 200 simulation runs. Vertical panels represent varying correlation between active clusters. Horizontal panels represent different signal-to-noise ratios.

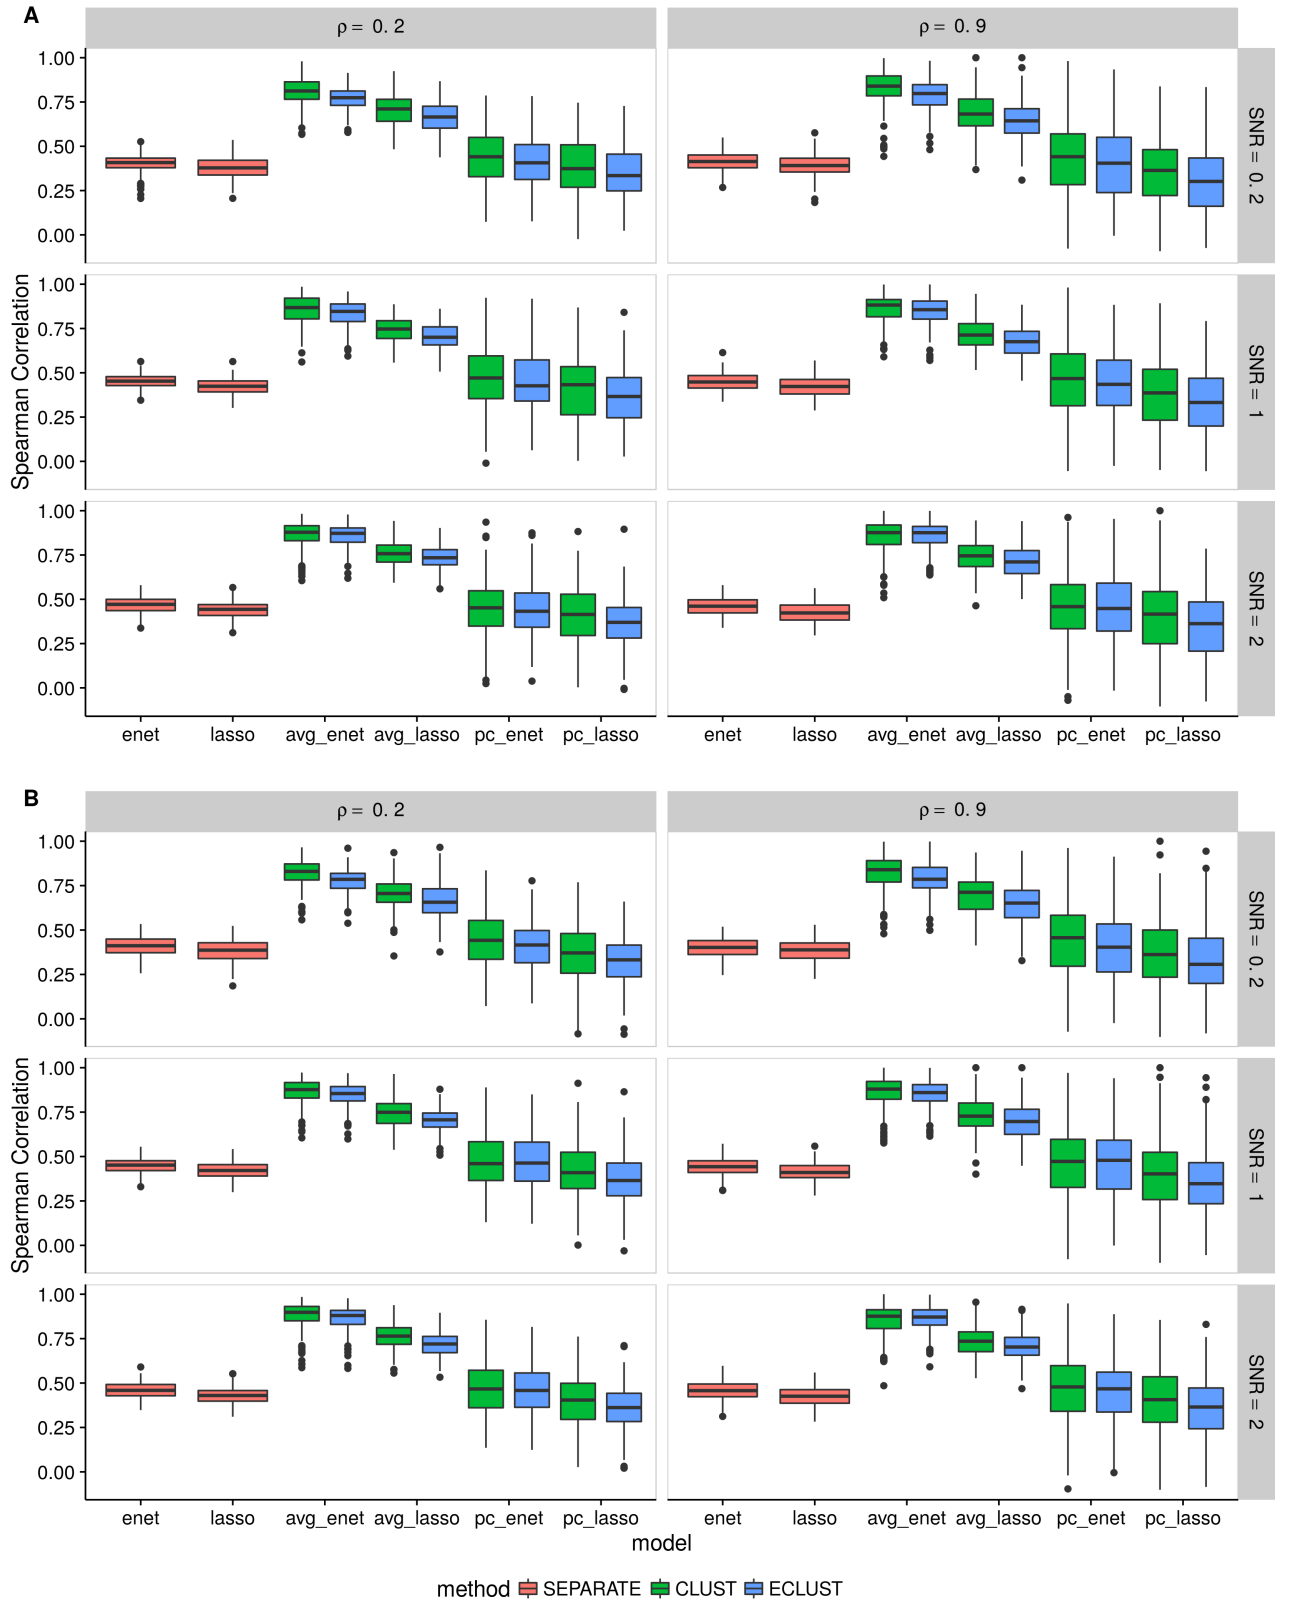


Figure S36: Simulation 2 – Average Spearman correlation from 10 CV folds of the training set using the Pearson correlation as a measure of similarity. (A) $\alpha_j \sim \text{Unif}[0.4, 0.6]$, (B) $\alpha_j \sim \text{Unif}[1.9, 2.1]$. We fit the model to each of the 10 CV folds resulting in 10 sets of estimated regression coefficients. We then calculate the Spearman correlation between all $\binom{10}{2}$ possible combinations of these sets and take the average. This process is repeated for each of the 200 simulation runs. Vertical panels represent varying correlation between active clusters. Horizontal panels represent different signal-to-noise ratios.

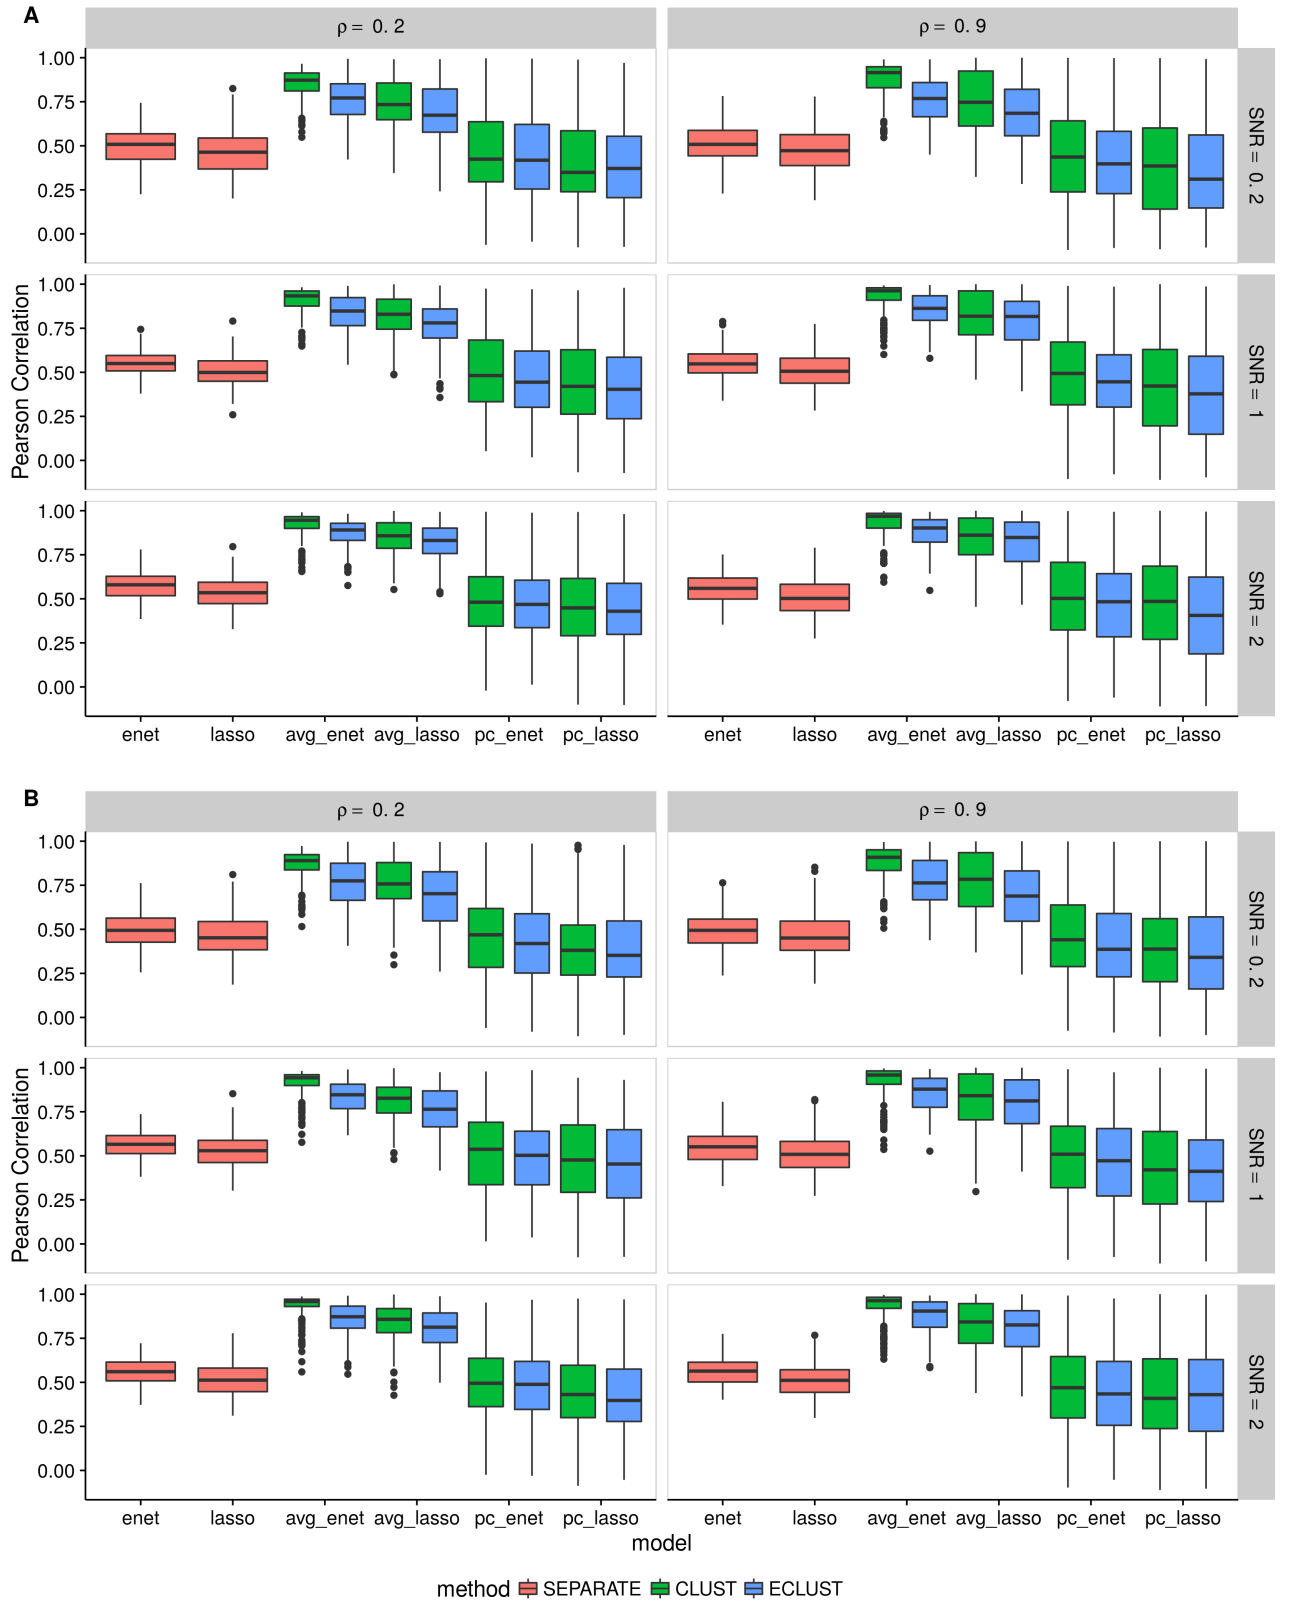


Figure S37: Simulation 2 – Average Pearson correlation from 10 CV folds of the training set using the Pearson correlation as a measure of similarity. (A) $\alpha_j \sim \text{Unif}[0.4, 0.6]$, (B) $\alpha_j \sim \text{Unif}[1.9, 2.1]$. We fit the model to each of the 10 CV folds resulting in 10 sets of estimated regression coefficients. We then calculate the Pearson correlation between all $\binom{10}{2}$ possible combinations of these sets and take the average. This process is repeated for each of the 200 simulation runs. Vertical panels represent varying correlation between active clusters. Horizontal panels represent different signal-to-noise ratios.

E.3 Simulation 3




Figure S38: Simulation 3 – Root mean squared error on an independent test set using the Pearson correlation as a measure of similarity from 200 simulation runs. Vertical panels represent varying correlation between active clusters. Horizontal panels represent different signal-to-noise ratios.




Figure S39: Simulation 3 – Correct Sparsity based on the training set using the Pearson correlation as a measure of similarity from 200 simulation runs. Vertical panels represent varying correlation between active clusters. Horizontal panels represent different signal-to-noise ratios.




Figure S40: Simulation 3 – True positive rate vs. false positive rate based on the training set using the Pearson correlation as a measure of similarity. Each point represents 1 simulation run (there are a total of 200 simulation runs). Vertical panels represent varying correlation between active clusters. Horizontal panels represent different signal-to-noise ratios.




Figure S41: Simulation 3 – Average Jaccard Index from 10 CV folds of the training set using the Pearson correlation as a measure of similarity. We fit the model to each of the 10 CV folds resulting in 10 sets of selected predictors. We then calculate the Jaccard Index between all $\binom{10}{2}$ possible combinations of these sets and take the average. This process is repeated for each of the 200 simulation runs. Vertical panels represent varying correlation between active clusters. Horizontal panels represent different signal-to-noise ratios.

F Visual Representation of Similarity Matrices

F.1 Pearson Correlation Matrix




Figure S42: Pearson correlation matrices of simulated predictors based on subjects with (a) $E = 0$, (b) $E = 1$, (c) their absolute difference and (d) all subjects. Dendrograms are from hierarchical clustering (average linkage) of one minus the correlation matrix for a, b, and d and the euclidean distance for c. The *module* annotation represents the true cluster membership for each predictor, and the *active* annotation represents the truly associated predictors with the response.

References

- Ravasz, E., Somera, A. L., Mongru, D. A., Oltvai, Z. N., and Barabási, A.-L. (2002). Hierarchical organization of modularity in metabolic networks. *science*, 297(5586):1551–1555.
- Zhang, B. and Horvath, S. (2005). A general framework for weighted gene co-expression network analysis. *Statistical applications in genetics and molecular biology*, 4(1).