THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Controlling false discoveries in Bayesian gene networks with lasso regression p-values

Citation for published version:

Wang, L & Michoel, T 2017 'Controlling false discoveries in Bayesian gene networks with lasso regression p-values' ArXiv.

Link:

[Link to publication record in Edinburgh Research Explorer](#)

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Subject Section

Controlling false discoveries in Bayesian gene networks with lasso regression p-values

Lingfei Wang* and Tom Michoel

Department of Genetics & Genomics, The Roslin Institute, The University of Edinburgh, Easter Bush, Midlothian EH25 9RG, UK.

*To whom correspondence should be addressed.

Associate Editor: XXXXXXX

Received on XXXXX; revised on XXXXX; accepted on XXXXX

Abstract

Motivation: Bayesian networks can represent directed gene regulations and therefore are favored over co-expression networks. However, hardly any Bayesian network study concerns the false discovery control (FDC) of network edges, leading to low accuracies due to systematic biases from inconsistent false discovery levels in the same study.

Results: We design four empirical tests to examine the FDC of Bayesian networks from three p -value based lasso regression variable selections — two existing and one we originate. Our method, `lassopv`, computes p -values for the critical regularization strength at which a predictor starts to contribute to lasso regression. Using null and Geuvadis datasets, we find that `lassopv` obtains optimal FDC in Bayesian gene networks, whilst existing methods have defective p -values. The FDC concept and tests extend to most network inference scenarios and will guide the design and improvement of new and existing methods. Our novel variable selection method with lasso regression also allows FDC on other datasets and questions, even beyond network inference and computational biology.

Availability: `Lassopv` is implemented in R and freely available at <https://github.com/lingfeiwang/lassopv> and <https://cran.r-project.org/package=lassopv>.

Contact: Lingfei.Wang@roslin.ed.ac.uk

Supplementary information: Supplementary data are available at *Bioinformatics* online.

1 Introduction

The reconstruction of gene regulation networks from gene expression data has been one of the major interests and challenges in computational and systems biology ([1, 2]). As opposed to co-expression networks, directed networks are specific on the source and target of gene regulations and, consequently, have attracted increasing attention. Among them, Bayesian networks model both the correlation structure among gene expression profiles and the conditional independencies between them, using directed acyclic graphs (DAGs), such as in [3, 4, 5, 6]. An accurate Bayesian network can significantly improve downstream analyses and experimental validations by predicting the master regulators and the genes responding to them in specific experimental conditions [5, 6, 7].

Despite such advantages of Bayesian networks over co-expression networks, its false discovery control (FDC), i.e. setting a network sparsity threshold such that the expected number of false positive edges is controlled, has been left mostly untouched as a statistical hypothesis testing

question. Co-expression networks contain a p -value for every pair of genes, from which a threshold for FDC can be chosen (e.g. [8]), despite the low specificity. However, no Bayesian network inference method to our knowledge provides a spectrum of continuous values (like p -value) for FDC. Instead, Bayesian network inference has been regarded as either a mathematical optimization problem or a multi-step, multi-parametric test [9, 10], so FDC has become very difficult. The importance of FDC, or even FDC itself, has been mostly overlooked in Bayesian networks.

The consequences of lacking FDC are obvious. Without a p -value-like variable to control the false discovery rate, any choice of network sparsity is hard to justify statistically. An even more detrimental consequence is the sub-optimal inference accuracy due to one group of interactions being favored over the rest by statistical bias, as discussed in detail in Section 2. Testing and achieving FDC would allow us to evaluate and improve network inference methods.

To control the false discovery, in this paper we consider the Bayesian network inference when a natural ordering of genes is given. In this case, Bayesian network inference reduces to a series of individual variable

selection problems [11]. Unlike traditional likelihood optimization or sampling algorithms, these approaches also easily scale to large systems involving thousands of genes. Furthermore, they are not hampered by the fact that DAGs form equivalence classes, in which individual DAGs cannot be distinguished by observational data alone [12]. Natural gene orderings or (dense) prior DAGs are typically obtained from external information (e.g. signalling pathways or known transcription factor binding target information)[13], or, in a systems genetics context, can be inferred using expression quantitative trait loci in the neighborhood of each gene’s transcription start site (cis-eQTL) as causal anchors [14, 15, 16]. By then applying regression methods, particularly the lasso [17, 18, 19] which favors sparse solutions, we can perform variable selection on the candidate regulators of each gene, and obtain sparse, genome-scale, and high-quality Bayesian networks [20, 13, 21].

There have been several attempts to repurpose lasso regression for statistical variable selection. Cross-validation, the standard approach in predictive lasso regression, is computationally intensive and disregards variable selection FDC [13]. Scaling regularization strengths with the number of candidate parental genes, as proposed in [20], offsets more parents with stronger selection, and can provide an upper bound for false discovery. Without a lower bound, however, its over-conservative FDC is still subject to biases dependent on the number of parental genes. None of these methods could demonstrate consistent FDC in network inference or variable selection.

Recent developments in sequential hypothesis testing regarded lasso regression as hypothesis testing in variable selection [22, 23, 24]. By considering every candidate predictor separately, the null hypothesis assumes an independent predictor, from which the p -value of specific lasso statistics of interest may be obtained. Accurate p -values of lasso regression variable selection would make possible FDC in Bayesian networks, like in co-expression networks and association studies. Consistent FDCs across all target genes would allow for optimal inference accuracy and a self-justified network sparsity.

In this paper, we develop the software `lassopv` to compute the p -value of the critical regularization strength at which every predictor starts to contribute to lasso regression for the first time [25]. This choice of statistic is based on the established criterion that predictors with nonzero contributions are significant, and more so at stronger regularizations. We also propose four statistical tests to empirically evaluate the FDC of a reconstructed Bayesian network. To compare `lassopv`’s FDC against existing methods (`covTest` [22] and `selectiveInference` [23]), we apply these statistical tests on their reconstructed networks from four low-dimensional and high-dimensional, simulated null and real biological datasets [26]. We also demonstrate Bayesian networks’ advantage over co-expression networks in avoiding indirect regulations and confounding. `Lassopv` is publicly available in R at <https://github.com/lingfeiwang/lassopv> and <https://cran.r-project.org/package=lassopv>.

2 Approach

The critical question in Bayesian network inference is orienting the regulations, which can be achieved with three broad categories of methods. The first, known as the PC algorithm [27, 9], tries to identify the V-shaped interactions $G_1 \rightarrow G_2 \leftarrow G_3$ between three genes G_1, G_2, G_3 , and then propagates the orientations on the co-expression network. The second method treats every Bayesian network as a regression model, and seeks the best predictive network(s) with Monte Carlo simulations [10]. The third method, as discussed before, introduces external information to order all genes, and then network inference becomes a series of regression problems

of every gene on its predecessors in the ordering. Hybrid methods have also been developed.

Regardless of the category, however, the form of their products is identical — a list of gene regulations that comprise the Bayesian network at any given significance level, regularization strength, or other custom cutoff value. This threshold determines network sparsity: starting with no interaction, and as the threshold changes monotonically, the network gradually includes more and more interactions, and finally reaches a full network (with $n(n-1)/2$ interactions for n genes). Therefore, every Bayesian network inference method, regardless of the category, consists effectively of two separate and sequential questions — first how to orient all regulations, answered by the full network, and then what value to assign to each regulation to rank their significance, answered by the sequence and/or critical threshold at which they become significant in the network.

In this paper, we focus on the FDC of the latter question of value assignment. To understand how the lack of FDC can undermine inference accuracy, as mentioned in Section 1, consider an example association study for multiple diseases whose sample sizes differ. If the correlation coefficient or odds ratio was used as a significance measure, the study could not account for the sample size differences, and therefore would bias towards associations with diseases with fewer samples. Such biases lead to inconsistent FDC levels between different groups of tests (here diseases), and therefore reduce the overall accuracy. The proper significance measure in the example would be the p -value (or false discovery rate) for observing a correlation coefficient or odds ratio value, which can ensure a unified FDC and an optimal accuracy across multiple tests. In network inference, statistical tests are composite, sequential, and much more complicated than simple pairwise correlations. Therefore, such biases cannot be assumed absent with merely equal sample sizes, but should be tested.

Testing FDC bias in network inference is equivalent to testing (pseudo-)random number generators in cryptography. Optimal FDC prevents any bias, and therefore all the non-existent candidate interactions (false cases) would appear identical to the network inference method, which can only assign values that rank them randomly and featurelessly. Although the possible feature space has infinite dimensions, there are typical bias features which network inference methods tend to produce, and therefore should be tested, as for random number generators [28]. For example, after orienting the interactions, the different numbers of candidate incoming interactions for different genes may introduce a bias, especially for regression based network inference.

However, most Bayesian network inferences only produce a list of significant interactions without specifying their values or rankings, posing a major difficulty to the empirical evaluations of FDC. Under the null hypothesis, each interaction has an equal probability of being significant, at any network sparsity. Therefore, for any target gene, the expected number of its significant incoming interactions (i.e. false positives) should be the product of that of its potential ones (i.e. false cases) and the false positive rate (FPR). As a result, given only a list of significant interactions, a consistent FPR would yield a testable linear relation between the numbers of significant and possible incoming interactions.

Besides the linearity test, we devise additional tests for Bayesian networks of continuous values or rankings of potential interactions, such as those obtained from lasso regression p -values. We also account for other issues in FDC on real data, such as the high dimensionality in genomic data and the “contamination” from unknown genuine interactions.

3 Methods

3.1 Data

We used the Geuvadis consortium’s gene expression levels and SNPs from lymphoblastoid cell lines of 360 European individuals [26] for gene network reconstruction. We limited our analyses to the 3172 genes that possess at least one significant cis-eQTL [26]. The expression levels of every gene were converted into following the standard normal distribution by relative ranking [14] prior to subsequent analyses.

We used the most significant cis-eQTL for every gene as the causal anchor [29, 16] to infer the probability of directed regulation between all pairs of the 3172 genes, using the function `pij_gassist` in `Findr 1.0.5` [16, 30]. Based on the inferred probability of pairwise regulation, we then constructed a greedy, maximal DAG by adding one directed regulation at a time in descending probability order using `netr_one_greedy` also in `Findr 1.0.5`. Edges that were to create a loop were discarded. The maximal DAG of 5,029,206 edges represented the natural prior ordering of the 3172 genes which, together with their expression levels across 360 individuals, formed the input of lasso-based Bayesian network learning.

Based on the above full, high-dimensional Geuvadis dataset, we additionally derived three datasets to evaluate lasso p -values for Bayesian network learning under different conditions:

- The high-dimensional null dataset of the same dimension (3172 genes from 360 individuals) was constructed by replacing every expression level with a random sample from independent standard normal distribution. This dataset examines the performance of variable selection methods in a high-dimensional null setting.
- The low-dimensional Geuvadis dataset of the top 150 genes in the prior order from all 360 individuals validates the low-dimensional real performance.
- The low-dimensional null dataset was similarly constructed from the low-dimensional Geuvadis dataset, by replacing every expression level with a random sample of independent standard normal distribution, to reflect the low-dimensional null scenario.

3.2 Lasso p -value

Without loss of generality, assume we have n observations of the target variable as $\mathbf{y} \in \mathbb{R}^n$ and k predictor variables as $\mathbf{X} \equiv (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_k) \in \mathbb{R}^{n \times k}$, subject to normalization:

$$\text{mean}(\mathbf{y}) = 0, \quad \text{and } \forall i = 1, \dots, k, \quad \text{mean}(\mathbf{x}_i) = 0. \quad (1)$$

Parameterized by the regularization strength λ , the lasso regularization [17] solves the optimization problem

$$\hat{\beta}(\lambda) \equiv \underset{\beta \in \mathbb{R}^k}{\text{argmin}} \frac{1}{2n} \|\mathbf{y} - \mathbf{X}\beta\|_2^2 + \lambda \|\beta\|_1, \quad (2)$$

and predicts $\mathbf{y}$ with the estimator $\hat{\mathbf{y}}(\lambda)$ as

$$\hat{\mathbf{y}}(\lambda) \equiv \mathbf{X}\hat{\beta}(\lambda). \quad (3)$$

To compute the p -value for predictor $\mathbf{x}_i$, first denote its real variance with $\sigma_i^2 \equiv \frac{1}{n} \|\mathbf{x}_i\|_2^2$. Its null hypothesis $H_0^{(i)}$ can be defined as a uniformly distributed vector on the sphere $\mathbb{S}^{n-1}$ with the same variance/radius, *i.e.*

$$P(\mathbf{x}_i = \mathbf{v} \mid H_0^{(i)}) = \begin{cases} \text{const}, & \|\mathbf{v}\|_2 = \sigma_i \sqrt{n}, \\ 0, & \text{else.} \end{cases} \quad (4)$$

We chose the test statistic as the critical regularization strength $\lambda = \lambda_i$ at which predictor $\mathbf{x}_i$ first contributes to predicting $\mathbf{y}$, as

$$\lambda_i \equiv \sup\{\lambda : \hat{\beta}_i(\lambda) \neq 0\}. \quad (5)$$

The one-sided p -value of the critical regularization strength $\tilde{\lambda}_i$ of $\mathbf{x}_i$ in the regularization path is the probability of $\lambda_i \geq \tilde{\lambda}_i$ under the null hypothesis, which can be analytically computed under approximation (Section S1) as

$$P(\lambda_i \geq \tilde{\lambda}_i \mid H_0^{(i)}) \approx 2 - 2\text{CDF}_{t(n-2)}\left(\frac{\tilde{\lambda}_i}{\sigma_i \sigma_{\mathbf{y}_{\text{res}}}(\tilde{\lambda}_i)} \sqrt{\frac{n-2}{1 - \frac{\tilde{\lambda}_i^2}{\sigma_i^2 \sigma_{\mathbf{y}_{\text{res}}}^2(\tilde{\lambda}_i)}}}\right), \quad (6)$$

where $\text{CDF}_{t(n-2)}$ represents the cumulative distribution function for the Student’s t -distribution with $n - 2$ degrees of freedom, and

$$\sigma_{\mathbf{y}_{\text{res}}}^2(\lambda) \equiv \frac{1}{n} \|\mathbf{y} - \hat{\mathbf{y}}(\lambda)\|_2^2 \quad (7)$$

is the residue variance of $\mathbf{y}$ at any given λ .

Besides the ‘spherical’ null hypothesis above, in Section S1 we also considered the ‘normal’ null hypothesis (Eq S5) of independent normal distributions with variance σ_i^2 . We developed the package `lassopv` [25] to compute the p -values for these null hypotheses.

3.3 False discovery control using p -values in network inference

To prune a prior DAG of possible regulations as set $\mathcal{N}_0 \equiv \{i \rightarrow j\}$ of n genes $i, j \in \{1, \dots, n\}$, we need to compute the p -value of every regulation $i \rightarrow j \in \mathcal{N}_0$ as p_{ij} , separately for each target node j as a variable selection. Noting that the null p -values follow the standard uniform distribution $p_{ij} \sim U(0, 1)$, we evaluated the quality of null p -values and their FDC with the following tests:

- **Histogram test:** To evaluate the null p -values as a whole, we visualized the overall histogram of p_{ij} against the standard uniform histogram of the same observation size.
- **The Kolmogorov-Smirnov (KS) test:** To evaluate the null p -values separately for each variable selection, we performed the KS test [31, 32] on p_{ij} against the standard uniform distribution (separately for each target gene j but over all its possible regulators i). A Manhattan plot then shows the KS test p -values as a function of the number of possible regulators $N_j \equiv \#\{i \rightarrow j \in \mathcal{N}_0\}$. The KS p -values were then compared against the 0.05 significance threshold with Bonferroni correction.
- **Linearity test:** To test whether p -values can be compared across different variable selections in network inference, we can choose a significance threshold p_{thres} for p -value. Under the null hypothesis, the number of significant regulators for every target gene i would approximately follow the normal distribution $N(N_i p_{\text{thres}}, N_i p_{\text{thres}}(1 - p_{\text{thres}}))$ (or strictly speaking the binomial distribution $B(N_i, p_{\text{thres}})$). For a given threshold p_{thres} , we can visually examine the linear relation between the numbers of possible and significant regulators for each target in a scatter plot.
- **R^2 test:** Based on the linearity test, we further computed the goodness of linear fit R^2 at different significance thresholds, and as a function of the total proportion of significant hypotheses. A higher curve or a larger area under the R^2 curve (AUR2) then indicates better FDC. Since the accuracy of small p -values are most important for distinguishing null and non-null cases, partial AUR2s at small p -values were also computed.

Although these tests were designed primarily for p -values, the linearity and R^2 tests also apply to other continuous statistics straightaway. For the histogram and KS tests, any other continuous statistic (or biased p -values) can also be compared against its (empirical) null distribution, if available.

Given only a list of significant edges, the linearity and R^2 tests are also test the FDC of the reconstructed binary Bayesian networks.

These tests also assume a null dataset, i.e. the absence of any genuine interactions in the network. To extend them to real datasets where sparse real interactions lead to the enrichment of small p -values or other statistics, a proper exclusion cutoff can remove those “contaminations”. The specific treatments are explained in the relevant results section.

3.4 Evaluation of existing and new lasso p -values

We limited our evaluation and comparison among the three R packages that compute lasso p -values:

- `lassopv` 0.2.0: this paper and [33],
- `covTest` 1.02: covariance test on the local loss of explained variance when a predictor is removed [22], and
- `selectiveInference` 1.2.0: post-selection inference based on sequential hypothesis testing (at multiple values of its parameter λ , [23, 24]).

Methods that cannot give p -values were not included, such as [13] (knockoff) and [34].

We evaluated each method on each dataset. Given the inferred maximal DAG, we performed one variable selection per target gene, by computing the lasso p -values of all its potential regulators. Using tests in Section 3.3, we then evaluated the FDC of p -values from different packages and different datasets.

4 Results

4.1 Lasso p obtained uniformly distributed p -values on the low-dimensional null dataset

The signalling pathways or gene regulatory networks of biological systems can be modeled as a sparse DAG or Bayesian network, inferred from observational gene expression data [35]. When a superset DAG or a natural ordering of (gene) nodes is given, the problem reduces to a series of variable selections, where each node is regressed on its parents (i.e. potential regulators) in the superset DAG [13, 21]. An accurate p -value based FDC would allow for a uniform significance threshold across the regressions, and therefore an optimal overall accuracy for network inference.

To evaluate the FDC of different lasso p -value methods in network inference, we first looked into the baby problem with low-dimensional null datasets and used `lassopv`, `covTest`, and `selectiveInference` to prune the maximal DAG into sparse DAGs. Considering all variable selections together in the histogram test in Figure 1 (histogram), `Lasso $p$` produced uniformly distributed p -values on the null dataset. `CovTest` was highly over-conservative, with over-abundance of p -values towards one. `SelectiveInference` had p -values that were over-abundantly small and large at high λ but were uniformly distributed at low λ , in agreement with Figures 3 and 5 in [23].

We then evaluated variable selections separately for each target gene using the KS test, and `lassopv` maintained its optimal FDC (Figure 1, KS). `Lasso $p$` p -values were consistent with the standard uniform distribution on the low-dimensional null dataset. Although `selectiveInference` obtained an overall standard uniform distribution in the histogram test, its null p -value distributions were revealed highly non-uniform for many target genes. For `covTest`, we accounted for its biased p -values by performing the KS test against its overall empirical p -value distribution. However, its null p -value distributions still differ across different target genes, suggesting that the amount of bias is dependent on the number of possible regulators.

To specifically evaluate the FDC in network inference, we then performed the R^2 test on the low-dimensional null dataset, which

Table 1. `Lasso $p$` outperformed `covTest` and `selectiveInference` on AUR2 (Figure 1) of the low-dimensional null dataset. Partial AUR2s were computed for the given x bounds, and normalized to unit AUR2 at constant function $R^2 = 1$. Best performers are shown in bold.

Bound	lassopv	covTest	selectiveInference	
			$\lambda = 0.001$	0.02
[0, 0.01]	0.12	0.032	0.010	0.008
[0, 0.05]	0.29	0.031	0.051	0.037
[0, 0.2]	0.63	0.21	0.22	0.19
[0, 1]	0.90	0.76	0.75	0.72

Table 2. `Lasso $p$` outperformed `covTest` and `selectiveInference` on AUR2 (Figure 2) of the low-dimensional Geuvadis dataset. The layout is the same with Table 1.

Bound	lassopv	covTest	selectiveInference	
			$\lambda = 0.001$	0.02
[0, 0.01]	0.049	0.003	0.013	0.007
[0, 0.05]	0.11	0.008	0.010	0.034
[0, 0.2]	0.35	0.094	0.11	0.15
[0, 1]	0.81	0.70	0.70	0.66

reaffirmed our existing conclusions (Figure 1, R^2). `Lasso $p$` obtained a highly linear relation between the numbers of candidate and significant regulators, especially at small p -values, and as opposed to `covTest` and `selectiveInference`. `SelectiveInference` assigned too many highly significant regulators to a small number of target genes, which could be the cause of breakdown in the KS test. Their performances were also summarized in full and partial AUR2s in Table 1.

4.2 Lasso p obtained optimal FDC on the low-dimensional Geuvadis dataset

To evaluate the FDC of different lasso p -value methods on the inference of a real gene regulation network, we first looked into the low-dimensional Geuvadis dataset ([26], Section 3.1). Despite our lack of knowledge on the groundtruth of the gene regulation network, the sparse nature of incoming regulations for every gene still makes the FDC evaluation possible using the same tests we applied on the null dataset.

Starting from the maximal DAG of the low-dimensional Geuvadis dataset, we first performed the overall histogram test on the results of `lassopv`, `covTest`, and `selectiveInference` (Figure 2, histogram). Compared to the low-dimensional null dataset (Figure 1), All three methods yielded an excess of small p -values, reflecting the existence of genuine gene interactions. However, for `selectiveInference`, the excess was scarce at low λ , and overlapped with false positives at high λ . Consequently, `selectiveInference` had a low specificity on the Geuvadis data, which could not be resolved with finer choices of λ (not shown). This also showed the difficulty in choosing the correct λ for `selectiveInference`.

On the per-target evaluation with KS test, we excluded the genuine interactions by discarding the p -values smaller than 0.01. The remaining interactions were then compared against the uniform distribution $U(0.01, 1)$. For `covTest`, we removed the bottom 1% of all p -values to account for its biased null distribution. Results were in agreement with those of the null dataset (Figure 2, KS), confirmed the ideal FDC from `lassopv`, and remained stable at higher exclusion thresholds (not shown).

The existence of genuine interactions violated null linearity and lowered R^2 for all three methods, more so on the leading method `lassopv` (Figure 2, R^2). This was also shown in the AUR2 of these methods (Table 2). As expected, the violation appeared small and localized

Fig. 1. `lassopv` provided accurate FDC than `covTest` and `selectiveInference` in the histogram, KS, and R^2 tests on low-dimensional null dataset. Histogram: the histograms of lasso p -values, with dashed lines as the perfectly uniform histogram of the same size. KS: the Manhattan plots of KS test p -values (y) of lasso p -values against the standard uniform distribution (or the overall empirical distribution for `covTest`) as a function of the numbers of potential regulators of each gene (x). Dashed lines represent the KS test p -value 0.05 with Bonferroni correction. High dots and stars (too small for machine precision) indicate p -values significantly deviating from the desired distribution. R^2 : the R^2 goodness of fit (blue, left y) and the maximum number of significant regulators for all genes (red, right y) as functions of the total proportion of significant regulations (x) at different significance thresholds. Higher blue and linear red curves indicate better FDC across the variable selection tasks. See Methods for test methods.

Fig. 2. `lassopv` provided accurate FDC than `covTest` and `selectiveInference` according to the histogram, KS, and R^2 tests on low-dimensional Geuvadis dataset. The layout is the same with Figure 1.

Fig. 3. Lassopv provided accurate FDC on the high-dimensional null dataset. (A, B, C) The histogram (A), KS (B), and R^2 (C) tests following the same layout as Figure 1. To account for high dimensionality, we removed p -values=1 in (A), and only tested the bottom 5% p -values in (B). (D, E, F) The linearity test for lassopv between the numbers of predictors (x) and significant predictors (y) at p -value significance thresholds of the bottom 1% (D, $R^2 = 0.84$), 5% (E, $R^2 = 0.96$), and 20% (F, $R^2 = 0.80$) potential regulations. Every dot corresponds to one variable selection in the network inference.

at small p -values and low recall, because the number of genuine regulators is low for every target gene. Other than that, method performances mostly agreed with those on the low-dimensional null dataset, including that selectiveInference continued to assign highly significant regulators to a small number of targets.

In summary, the performances of all three methods on the low-dimensional Geuvadis dataset agreed highly with those on the low-dimensional null dataset. This has several implications. First, the statistical tests of FDC in Section 3.3 could also be applied on real gene regulation networks, after adjustments for sparse interactions. Second, the null dataset was validated to highly resemble the null interactions in the Geuvadis dataset. This supported our upcoming FDC evaluations with the simulated high-dimensional null dataset. Third, we continued to find lassopv as the best method for FDC.

4.3 Lassopv obtained accurate false discovery control on the high-dimensional datasets

In high-dimensional problems, lasso regression becomes under-determined as the regularization strength approaches zero. Machine precision prevents some of the insignificant regulators from having any role in the regression, and therefore biases their p -values to one. This challenge can only be overcome by a high-precision lasso solver. However, this would not affect our analysis of identifying significant regulations, because we are only interested in the small p -values.

For that reason, we tested lassopv and covTest directly on the high-dimensional null dataset containing 3172 genes. SelectiveInference could not handle high-dimensional scenarios and therefore was excluded

Table 3. Lassopv outperformed covTest and selectiveInference on AUR2 (Figure 3C, Figure 4, Figure 5) of the high-dimensional null and Geuvadis datasets. Best performers are shown in bold.

Bound	Null dataset		Geuvadis dataset	
	lassopv	covTest	lassopv	covTest
[0, 0.01]	0.65	0.008	0.01	0.06
[0, 0.05]	0.88	0.10	0.31	0.06
[0, 0.2]	0.94	0.08	0.75	0.06

Fig. 4. covTest failed to provide FDC on the high-dimensional null dataset. The layout is the same with Figure 3. (B) the KS test was performed against the empirical distribution of all p -values (of the bottom 5%). (D, E) covTest could not provide FDC in network inference by failing the linearity test, at p -value significance thresholds of the bottom 1% (D, $R^2 = 0.03$) and 5% (E, $R^2 = 0.13$). (F) covTest claimed all gene regulations significant at the bottom 20% threshold due to ties.

from the analysis. As confirmed for lassopv in Figure 3A, the high-dimensional effect unavoidably biased a large proportion of insignificant regulations to p -value=1 (removed in figure). The resulting p -value distribution differed notably from the standard uniform distribution, but only for insignificant regulations (p -value>0.1). In spite of that, its FDC remained accurate because p -values<0.1 were still uniformly distributed. The slight over-abundance of p -values between 0.1 and 0.2 was due partly to the same high-dimensional effect, and partly to the analytical approximation in Section S1. A simulation-based p -value computation without the approximation shrank the peak by half (not shown).

Lassopv also displayed accurate FDC in other tests. The KS test on the bottom 5% p -values confirmed lassopv’s accurate FDC on a per variable selection basis (Figure 3B). In network inference, lassopv attained a highly linear relation ($R^2 = 0.96$) between the numbers of potential and significant regulators (Figure 3DE), as a typical example of what a FDC conforming network reconstruction method should yield in the linearity test. Although the high-dimensional effect resulted in plateaus at weak thresholds (Figure 3CF), these again did not affect FDC at small p -values. Its proper network FDC was also demonstrated in the R^2 curve (Figure 3C) and AUR2 (Table 3).

Fig. 5. *lassopv* provided accurate FDC than *covTest* on the high-dimensional Geuvadis dataset. The layout is the same with Figure 1. Besides, to remove the effects of high dimensionality and genuine interactions, only the bottom 3% to 5% (totalling 2%) of all p -values were evaluated in the KS tests.

On the other hand, *covTest* could not achieve FDC in high-dimensional network inference. Particularly, the nonlinear relation between the numbers of potential and significant regulators indicated a lack of FDC (Figure 4DE). This was also manifested in its failed KS and R^2 tests (Figure 4BC, Table 3). Besides that, the high-dimensional effect biased *covTest*’s insignificant regulations towards p -value=1 even more strongly than *lassopv*’s (Figure 4AF). The same conclusions hold for *lassopv* and *covTest* with the high-dimensional Geuvadis dataset. After accounting for the genuine interactions, *lassopv* again showed optimal FDC (Figure 5, Table 3).

Although the high dimensionality of modern datasets introduces unavoidable distortions in insignificant p -values, and the groundtruth is unavailable, we found that *lassopv* could still provide ideal FDC for small p -values, and therefore is unaffected by high-dimensional effects. We established the principles and tests to evaluate FDC in network inference through *lassopv*’s symbolic linearity between the numbers of candidate and significant regulators. *Lassopv* remained the optimal FDC method for high-dimensional network inference.

4.4 *Lassopv* reduced false discovery by removing the spurious indirect regulations in Pearson co-expression network

Hardly any software package on the market is capable of FDC in a high-dimensional network reconstruction problem. To our knowledge, the only available ones are co-expression networks, except the *covTest* which is already shown highly biased. However, co-expression networks are well known for their high FPRs from indirectly regulated and confounded genes, due to its pairwise nature. Regression based methods can reduce these false positives by accounting for the direct and indirect effects together.

We tested the reduction of false positives from *lassopv*, by comparing it with Pearson co-expression network on the high-dimensional Geuvadis dataset with 3172 genes. Using the same significance threshold

Fig. 6. *Lassopv* reduced the false positives in co-expression networks. (A) Venn diagram of regulation overlap. (B-E) Hypergeometric enrichment p -values (BD) and the total numbers (CE) of significant co-expressions (y) in the *lassopv* networks of indirect regulations (BC) and indirect regulations + confounding relations (DE) of a given step, at different significance thresholds of Bonferroni adjusted p -value (x). (F) Color legend for B-E. Colors represent different numbers of direct steps in indirect or confounding relations, where one step is the inferred network of direct regulations.

of 0.05 with Bonferroni corrections across all interactions, *lassopv* discovered 8091 regulations as opposed to 25088 in the co-expression network, which mostly overlap (Figure 6A). We derived indirect regulations of the *lassopv* network by connecting steps of significant direct regulations. Indirect regulations up to 4-step far were enriched with significant co-expressions (Figure 6B), and covered over 70% of all co-expressions at the adjusted p -value cutoff 0.05 (Figure 6C). Further inclusion of confounded genes led to a much denser network, with an enrichment of co-expressions only up to 3-step far (Figure 6D), but covering around 95% of all co-expressions (Figure 6E).

Although a high-quality gold standard is not available on real data, we still found convincing evidence supporting the reduction of indirect and confounding false positives by Bayesian networks, particularly *lassopv*, from co-expression networks.

5 Conclusion

In this article, we presented the problem of false discovery control in Bayesian network reconstruction and its impact on inference accuracy and network sparsity. We designed four statistical tests to evaluate the FDC of reconstructed networks on real genomic data. We proposed a new method to repurpose lasso regression for variable selection and computed its p -value in the program *lassopv* for consistent FDC in network inference. On simulated and real, low- and high-dimensional genomic datasets, our statistical tests revealed hidden defects in *covTest* and *selectiveInference*—the existing variable selection methods also

from lasso regression — and demonstrated the advantageous, accurate p -values and FDC from `lassopv`.

Already known in existing GWAS's, consistent FDC is crucial for optimal prediction accuracy. Consequently, FDC evaluations can advise and guide the application and design of network inference methods, and should be performed on all reconstructed Bayesian networks of any continuous value or merely a list of interactions, although this paper only evaluated different p -values. On the other hand, as a hypothesis testing method for variable selection, `lassopv` is potentially applicable, for example, on causal (expression) quantitative trait loci discovery and supervised dimensional reduction, and also on other disciplines beyond computational biology.

This paper is not concerned with how to obtain a prior ordering of gene nodes. We used genotype and gene expression data of the same individuals in population-based studies to orient the causal direction between any pair of genes [36, 37] and obtained a dense prior DAG. Depending on data availability, the prior ordering may also be obtained from known regulatory gene function annotations or regulatory interactions (e.g. protein-DNA interactions) [13], or simply from a greedy maximum-likelihood optimization in the absence of any additional information [20].

Similarly, this paper only reconstructs the structure of the Bayesian network. Since the Bayesian network reflects the conditional dependency between gene expression levels, a predictive model may be derived with an additional regression of any target gene on its parental genes in the reconstructed network. Under a linear approximation for the sparse regulatory system, an unpenalized linear regression may suffice.

Funding

The authors gratefully acknowledge the support by BBSRC (grant numbers BB/P013732/1 and BB/M020053/1).

References

- [1]Eric E. Schadt. Molecular networks as sensors and drivers of common human diseases. *Nature*, 461(7261):218–223, September 2009.
- [2]Mete Civelek and Aldons J. Lusis. Systems genetics approaches to understand complex traits. *Nature Reviews Genetics*, 15(1):34–48, January 2014.
- [3]Nir Friedman, Iftach Nachman, and Dana Pe'er. Learning Bayesian Network Structure from Massive Datasets: The Sparse Candidate Algorithm. In *Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence*, UAI'99, pages 206–215, San Francisco, CA, USA, 1999. Morgan Kaufmann Publishers Inc.
- [4]Nir Friedman, Michal Linial, Iftach Nachman, and Dana Pe'er. Using Bayesian Networks to Analyze Expression Data. *Journal of Computational Biology*, 7(3-4):601–620, August 2000.
- [5]Jun Zhu, Bin Zhang, Erin N. Smith, Becky Drees, Rachel B. Brem, Leonid Kruglyak, Roger E. Bumgarner, and Eric E. Schadt. Integrating large-scale functional genomic data to dissect the complexity of yeast regulatory networks. *Nature Genetics*, 40(7):854–861, July 2008.
- [6]Bin Zhang, Chris Gaiteri, Liviu-Gabriel Bodea, Zhi Wang, Joshua McElwee, Alexei A. Podtelezhnikov, Chunsheng Zhang, Tao Xie, Linh Tran, Radu Dobrin, Eugene Fluder, Bruce Clurman, Stacey Melquist, Manikandan Narayanan, Christine Suver, Hardik Shah, Milind Mahajan, Tammy Gillis, Jayalakshmi Mysore, Marcy E. MacDonald, John R. Lamb, David A. Bennett, Cliona Molony, David J. Stone, Vilmondur Gudnason, Amanda J. Myers, Eric E. Schadt, Harald Neumann, Jun Zhu, and Valur Emilsson. Integrated Systems Approach Identifies Genetic Nodes and Networks in Late-Onset Alzheimer's Disease. *Cell*, 153(3):707–720, April 2013.
- [7]H. Talukdar, H Foroughi Asl, R. Jain, R. Ermel, A. Ruusalepp, O. Franzén, B. Kidd, B. Readhead, C. Giannarelli, T. Ivert, J. Dudley, M. Civelek, A. Lusis, E. Schadt, J. Skogsberg, T. Michoel, and J.L.M Björkegren. Cross-tissue regulatory gene networks in coronary artery disease. *Cell Systems*, 2:196–208, 2016.
- [8]Peter Langfelder and Steve Horvath. WGCNA: an R package for weighted correlation network analysis. *BMC Bioinformatics*, 9:559, December 2008.
- [9]Markus Kalisch and Peter Bühlmann. Estimating High-Dimensional Directed Acyclic Graphs with the PC-Algorithm. *Journal of Machine Learning Research*, 8(Mar):613–636, 2007.
- [10]Marco Scutari. Learning Bayesian Networks with the bnlearn R Package. *Journal of Statistical Software*, 35(3), 2010.
- [11]D. Koller and N. Friedman. *Probabilistic Graphical Models: Principles and Techniques*. MIT Press, 2009.
- [12]David Maxwell Chickering. Learning Equivalence Classes of Bayesian-Network Structures. *Journal of Machine Learning Research*, 2(Feb):445–498, 2002.
- [13]Ali Shojaie and George Michailidis. Penalized likelihood methods for estimation of sparse high-dimensional directed acyclic graphs. *Biometrika*, 97(3):519–538, 2010.
- [14]Lin S. Chen, Frank Emmert-Streib, and John D. Storey. Harnessing naturally randomized transcription to infer regulatory relationships among genes. *Genome Biology*, 8:R219, 2007.
- [15]Joshua Millstein, Bin Zhang, Jun Zhu, and Eric E. Schadt. Disentangling molecular relationships with a causal inference test. *BMC Genetics*, 10:23, 2009.
- [16]Lingfei Wang and Tom Michoel. Efficient and accurate causal inference with hidden confounders from genome-transcriptome variation data. *PLOS Computational Biology*, 13(8):e1005703, August 2017.
- [17]Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, 58(1):267–288, 1996.
- [18]Bradley Efron, Trevor Hastie, Iain Johnstone, and Robert Tibshirani. Least angle regression. *The Annals of statistics*, 32(2):407–499, 2004.
- [19]Nicolai Meinshausen and Peter Bühlmann. High-dimensional graphs and variable selection with the lasso. *The Annals of Statistics*, pages 1436–1462, 2006.
- [20]Mark Schmidt, Alexandru Niculescu-Mizil, Kevin Murphy, et al. Learning graphical model structure using l1-regularization paths. In *AAAI*, volume 7, pages 1278–1283, 2007.
- [21]Bryon Aragam and Qing Zhou. Concave penalized estimation of sparse Bayesian networks. *arXiv preprint arXiv:1401.0852*, 2014.
- [22]Richard Lockhart, Jonathan Taylor, Ryan J. Tibshirani, and Robert Tibshirani. A significance test for the lasso. *Ann. Statist.*, 42(2):413–468, 04 2014.
- [23]Jonathan Taylor and Robert Tibshirani. Post-selection inference for l1-penalized likelihood models. *arXiv preprint arXiv:1602.07358*, 2016.
- [24]Jason D. Lee, Dennis L. Sun, Yuekai Sun, and Jonathan E. Taylor. Exact post-selection inference, with application to the lasso. *Ann. Statist.*, 44(3):907–927, 06 2016.
- [25]Lingfei Wang. `Lassopv`, 2017. doi: 10.5281/zenodo.594121.
- [26]Tuuli Lappalainen, Michael Sammeth, Marc R. Friedlander, Peter A. C. /'t Hoen, Jean Monlong, Manuel A. Rivas, Mar Gonzalez-Porta, Natalja Kurbatova, Thasso Griebel, Pedro G. Ferreira, Matthias Barann, Thomas Wieland, Liliana Greger, Maarten van Iterson, Jonas Almlöf, Paolo Ribeca, Irina Pulyakhina, Daniela Esser, Thomas Giger, Andrew Tikhonov, Marc Sultan, Gabrielle Bertier, Daniel G. MacArthur, Monkol Lek, Esther Lizano, Henk P. J. Buermans, Ismael Padiou, Thomas Schwarzmayr, Olof Karlberg, Halit Ongen, Helena Kilpinen, Sergi Beltran, Marta Gut, Katja Kahlem, Vyacheslav Amstislavskiy, Oliver Stegle, Matti Pirinen, Stephen B. Montgomery, Peter Donnelly, Mark I. McCarthy, Paul Flícek, Tim M. Strom, The Geuvadis Consortium, Hans Lehrach, Stefan Schreiber, Ralf Sudbrak, Angel Carracedo, Stylianos E. Antonarakis, Robert Hasler, Ann-Christine Syvanen, Gert-Jan van Ommen, Alvis Brazma, Thomas Meitinger, Philip Rosenstiel, Roderic Guigo, Ivo G. Gut, Xavier Estivill, and Emmanouil T. Dermitzakis. Transcriptome and genome sequencing uncovers functional variation in humans. *Nature*, 501(7468):506–511, 09 2013.
- [27]Joseph Ramsey, Peter Spirtes, and Jiji Zhang. Adjacency-faithfulness and Conservative Causal Inference. In *Proceedings of the Twenty-Second Conference on Uncertainty in Artificial Intelligence*, UAI'06, pages 401–408, Arlington, Virginia, United States, 2006. AUAI Press.
- [28]Lawrence Bassham, Andrew Rukhin, Juan Soto, James Nechvatal, Miles Smid, Elaine Barker, Stefan Leigh, Mark Levenson, Mark Vangel, David Banks, N. Heckert, and James Dray. A Statistical Test Suite for Random and Pseudorandom Number Generators for Cryptographic Applications, April 2010. doi: 10.6028/NIST.SP.800-22r1a.
- [29]Eric E Schadt, John Lamb, Xia Yang, Jun Zhu, Steve Edwards, Debraj GuhaThakurta, Solveig K Sieberts, Stephanie Monks, Marc Reitman, Chunsheng Zhang, et al. An integrative genomics approach to infer causal associations between gene expression and disease. *Nature Genetics*, 37(7):710–717, 2005.
- [30]Lingfei Wang. Findr 1.0.5 (Fast Inference of Networks from Directed Regulations), December 2017. doi: 10.5281/zenodo.1095208.
- [31]Andrej N Kolmogorov. *Sulla determinazione empirica di una legge di distribuzione*. 1933.

- [32]N. Smirnov. Table for Estimating the Goodness of Fit of Empirical Distributions. *The Annals of Mathematical Statistics*, 19(2):279–281, 1948.
- [33]Lingfei Wang. Lassopv 0.2.0, February 2018. doi: 10.5281/zenodo.1183245.
- [34]Rina Foygel Barber and Emmanuel J. Candès. Controlling the false discovery rate via knockoffs. *Ann. Statist.*, 43(5):2055–2085, 10 2015.
- [35]Nir Friedman. Inferring cellular networks using probabilistic graphical models. *Science*, 303(5659):799–805, 2004.
- [36]Matthew V Rockman. Reverse engineering the genotype–phenotype map with natural genetic variation. *Nature*, 456(7223):738–744, 2008.
- [37]Yang Li, Bruno M Tesson, Gary A Churchill, and Ritsert C Jansen. Critical reasoning on causal inference in genome-wide linkage and association studies. *Trends in Genetics*, 26(12):493–498, 2010.

Supplementary Information

S1 Computation of lasso p -values

To compute the null distribution for variable $\mathbf{x}_i$, we simply assume the null hypothesis $H_0^{(i)}$ holds and omit its conditioning. The p -value expression in Eq 6 can be decomposed into:

$$\begin{aligned} & P(\lambda_i \geq \tilde{\lambda}_i) \\ &= P(\sup\{\lambda : \hat{\beta}_i(\lambda) \neq 0\} \geq \tilde{\lambda}_i) \\ &= P(\hat{\beta}_i(\tilde{\lambda}_i) \neq 0) \\ &+ P(\sup\{\lambda : \hat{\beta}_i(\lambda) \neq 0\} \geq \tilde{\lambda}_i \wedge \hat{\beta}_i(\tilde{\lambda}_i) = 0). \quad (\text{S1}) \end{aligned}$$

The first term is the probability that $\mathbf{x}_i$ is *active* (i.e. $\hat{\beta}_i \neq 0$ by definition) at $\lambda = \tilde{\lambda}_i$, whilst the second is the probability that $\mathbf{x}_i$ is active at some $\lambda > \tilde{\lambda}_i$ but inactive at $\lambda = \tilde{\lambda}_i$.

For every predictor, starting and stopping to be active are also called *entering and leaving the active set*, whose critical λ values are named *knots*. In lasso regression, predictors can enter and leave the active set multiple times. Therefore the second term in Eq S1 is nonzero in general. Here we make the approximation of neglecting the second term in Eq S1, which under-estimates p -values. However, we claim the consequent error is small, as justified in theory in Remark 1 and with data in Section 4.

Remark 1. Justification for neglecting the second term in Eq S1: Discovery-aimed variable selection focuses on selecting significant predictor variables, which should have small enough p -values. The second term in Eq S1 starts at zero and grows as more knots are gone through with decreasing λ . Significant variables become active earlier, and therefore this approximation incurs much smaller errors on them than in average, especially when the number of genuine predictors is small.

Remembering the variance of original $\mathbf{x}_i$ is σ_i^2 , here we define the variances of $\mathbf{x}_i$ following the null hypothesis as $\tilde{\sigma}_i^2$. Propositions are proven in Section S2, such as:

Proposition 1.

$$\hat{\beta}_i(\lambda) \neq 0 \iff \lambda < \tilde{\sigma}_i \sigma_{\mathbf{y}_{\text{res}}}(\lambda) |\text{Cor}(\mathbf{x}_i, \mathbf{y}_{\text{res}}(\lambda))|. \quad (\text{S2})$$

We then consider two possible null hypotheses separately.

S1.1 Normal null hypothesis

Here, assume the null hypothesis for $\mathbf{x}_i = (x_{1,i}, x_{2,i}, \dots, x_{n,i})^T$ is

$$H_0^{(i)} : x_{j,i} \sim i.i.d N(0, \sigma_i^2), \quad (\text{S3})$$

The following proposition is proven in Section S2:

Proposition 2. For $\mathbf{x} \in \mathbb{R}^n$ whose elements follow the i.i.d standard normal distribution, its dot product squared with another vector $\mathbf{y} \in \mathbb{R}^n$ follows the χ^2 distribution

$$n\sigma_{\mathbf{x}}^2 \text{Cor}^2(\mathbf{x}, \mathbf{y}) \sim \chi^2(1), \quad (\text{S4})$$

where $\sigma_{\mathbf{x}}^2$ is the variance of $\mathbf{x}$.

Therefore we obtain the p -value for every predictor i as

$$\begin{aligned} & P(\sup\{\lambda : \hat{\beta}_i(\lambda) \neq 0\} \geq \tilde{\lambda}_i) \\ &\approx P(\hat{\beta}_i(\tilde{\lambda}_i) \neq 0) \\ &= P(\tilde{\lambda}_i < \tilde{\sigma}_i \sigma_{\mathbf{y}_{\text{res}}}(\tilde{\lambda}_i) |\text{Cor}(\mathbf{x}_i, \mathbf{y}_{\text{res}}(\tilde{\lambda}_i))|) \\ &= P\left(n \left(\frac{\tilde{\sigma}_i}{\sigma_i}\right)^2 \text{Cor}^2\left(\frac{\mathbf{x}_i}{\sigma_i}, \mathbf{y}_{\text{res}}(\tilde{\lambda}_i)\right) > \frac{n\tilde{\lambda}_i^2}{\sigma_i^2 \sigma_{\mathbf{y}_{\text{res}}}^2(\tilde{\lambda}_i)}\right) \\ &= 1 - \text{CDF}_{\chi^2(1)}\left(\frac{n\tilde{\lambda}_i^2}{\sigma_i^2 \sigma_{\mathbf{y}_{\text{res}}}^2(\tilde{\lambda}_i)}\right), \quad (\text{S5}) \end{aligned}$$

where $\text{CDF}_{\chi^2(1)}$ represents the cumulative distribution function for distribution $\chi^2(1)$.

S1.2 Spherical null hypothesis

Now consider a restrictive null hypothesis, Eq 4, which fixes the variance of the null predictor ($\tilde{\sigma}_i = \sigma_i$). Then we have

$$\begin{aligned} & P(\sup\{\lambda : \hat{\beta}_i(\lambda) \neq 0\} \geq \tilde{\lambda}_i) \\ &\approx P(\hat{\beta}_i(\tilde{\lambda}_i) \neq 0) \\ &= P\left(|\text{Cor}(\mathbf{x}_i, \mathbf{y}_{\text{res}}(\tilde{\lambda}_i))| > \frac{\tilde{\lambda}_i}{\sigma_i \sigma_{\mathbf{y}_{\text{res}}}(\tilde{\lambda}_i)}\right) \\ &= 2 - 2\text{CDF}_{t(n-2)}\left(\frac{\tilde{\lambda}_i}{\sigma_i \sigma_{\mathbf{y}_{\text{res}}}(\tilde{\lambda}_i)} \sqrt{\frac{n-2}{1 - \frac{\tilde{\lambda}_i^2}{\sigma_i^2 \sigma_{\mathbf{y}_{\text{res}}}^2(\tilde{\lambda}_i)}}}\right) \quad (\text{S6}) \end{aligned}$$

where $t(n-2)$ is the Student's t -distribution with $n-2$ degrees of freedom.

Remark 2. Since $\lim_{n \rightarrow \infty} \tilde{\sigma}_i = \sigma_i$ in the normal null hypothesis, the spherical and normal null hypotheses (and their p -values) converge at large n .

Remark 3. The p -value of the first predictor under the spherical null hypothesis is the same with that of its Pearson correlation with the target vector $\mathbf{y}$.

S2 Proofs

S2.1 Proposition 1

First prove

$$\hat{\beta}_i(\lambda) > 0 \iff \lambda < \tilde{\sigma}_i \sigma_{\mathbf{y}_{\text{res}}}(\lambda) \text{Cor}(\mathbf{x}_i, \mathbf{y}_{\text{res}}(\lambda)). \quad (\text{S7})$$

Define

$$\tilde{\beta}_j(\lambda) \equiv \begin{cases} \alpha, & \text{for } j = i, \\ \hat{\beta}_j(\lambda), & \text{else.} \end{cases} \quad (\text{S8})$$

Since

$$\hat{\beta}_i(\lambda) \equiv \arg\min_{\alpha} \frac{1}{2n} \|\mathbf{y} - \mathbf{X}\tilde{\beta}\|_2^2 + \lambda \|\tilde{\beta}\|_1, \quad (\text{S9})$$

then

$$\begin{aligned} & \hat{\beta}_i(\lambda) > 0 \\ &\iff \frac{\partial}{\partial \alpha} \left(\frac{1}{2n} \|\mathbf{y} - \mathbf{X}\tilde{\beta}(\lambda)\|_2^2 + \lambda \|\tilde{\beta}(\lambda)\|_1 \right) \Big|_{\alpha=0+} < 0 \\ &\iff \frac{1}{n} \mathbf{x}_i^T (\alpha \mathbf{x}_i - \mathbf{y}_{\text{res}}(\lambda)) + \lambda \Big|_{\alpha=0+} < 0 \\ &\iff \lambda < \tilde{\sigma}_i \sigma_{\mathbf{y}_{\text{res}}}(\lambda) \text{Cor}(\mathbf{x}_i, \mathbf{y}_{\text{res}}(\lambda)). \quad (\text{S10}) \end{aligned}$$

Similarly,

$$\hat{\beta}_i(\lambda) < 0 \iff \lambda < -\tilde{\sigma}_i \sigma_{\mathbf{y}_{\text{res}}}(\lambda) \text{Cor}(\mathbf{x}_i, \mathbf{y}_{\text{res}}(\lambda)). \quad (\text{S11})$$

Since $\hat{\beta}_i(\lambda)$ and $\text{Cor}(\mathbf{x}_i, \mathbf{y}_{\text{res}}(\lambda))$ always have the same sign, we can combine the two as

$$\hat{\beta}_i(\lambda) \neq 0 \iff \lambda < \tilde{\sigma}_i \sigma_{\mathbf{y}_{\text{res}}}(\lambda) |\text{Cor}(\mathbf{x}_i, \mathbf{y}_{\text{res}}(\lambda))|. \quad (\text{S12})$$

S2.2 Proposition 2

Due to the rotational $SO(n)$ symmetry in the PDF of $\mathbf{x} \in \mathbb{R}^n$, each of which follows i.i.d $N(0, 1)$, the distribution of $n\sigma_{\mathbf{x}}^2 \text{Cor}^2(\mathbf{x}, \mathbf{y})$ does not depend on $\mathbf{y}$. For simplicity, choose

$$y_1 = \sqrt{n/2}, \quad y_2 = -\sqrt{n/2}, \quad y_i = 0, \text{ for } i = 3, \dots, n. \quad (\text{S13})$$

Then, by expanding the correlation, we have

$$n\sigma_{\mathbf{x}}^2 \text{Cor}^2(\mathbf{x}, \mathbf{y}) = (x_1/\sqrt{2} - x_2/\sqrt{2})^2. \quad (\text{S14})$$

Since $x_1, x_2 \sim i.i.d N(0, 1)$, define

$$z \equiv x_1/\sqrt{2} - x_2/\sqrt{2} \sim N(0, 1). \quad (\text{S15})$$

Therefore

$$n\sigma_{\mathbf{x}}^2 \text{Cor}^2(\mathbf{x}, \mathbf{y}) = z^2 \sim \chi^2(1). \quad (\text{S16})$$