Statistical power of gene-set enrichment analysis is a function of gene set correlation structure

DAVID M. SWANSON*

Oslo Centre for Biostatistics and Epidemiology, Oslo University Hospital, Oslo, NO 0407
dms866@mail.harvard.edu

Abstract

Motivation: We describe why statistical power for both self-contained and competitive gene-set tests is a function of the correlation structure of co-expressed genes, and why this characteristic is undesirable for gene-set analyses. Variable statistical power as a function of gene correlation structure has not been observed or studied previously. The observation is important in part because gene-set testing methodology is well-developed, yet this fundamental feature of many of its tests is unknown and has the potential to reinterpret past gene-set test results and guide future implementations, including those using sequence data. Type 1 error inflation is also amenable for study in our statistical framework; while it has been well-studied and described previously for both self-contained and competitive tests, it has less often been done in an analytical framework. Our observations apply to four commonly-used gene-set testing approaches for microarrays, including CAMERA, ROAST, SAFE, and GAGE, and a recently proposed one for RNAseq, MAST.

Results: We characterize situations in which power is especially small relative to effect sizes of genes in a set for both competitive and self-contained gene-set tests. We propose three alternative tests, one of which replicates the properties of permutation-based self-contained tests, but avoids the need for even recently proposed, rotation-based approximations to permutations. The two other proposed tests have the unique property that statistical power is not a function of co-expression correlation in the gene-set and therefore is the preferred methodology. We compare our proposed tests to leading gene-set tests and apply them to an already-published study of smoking exposure on pregnant women.

Contact: dms866@mail.harvard.edu

Supplementary Material: Online supplementary material includes additional simulation results supporting the relationship between the “mixed” and “directional” gene-set tests of ROAST and closed-form implementations of them.

1 Introduction

Gene set enrichment analysis (GSEA) or gene-set testing is a class of statistical methods whose goal is to assess the joint enrichment of a biologically interpretable set of genes in a microarray or RNAseq experiment^{1–6}. A variety of methods have been proposed nearly since the advent of microarrays, all with unique advantages, and there have been more recent modifications of these methods as RNAseq becomes a more common platform⁷. There have been two primary categories of gene-set tests, “competitive” and “self-contained”^{8–11}. The former tries to answer whether genes in a set are more differentially expressed than some background level of differential expression on the array. The latter tries to answer whether the gene-set is more differentially expressed than one were to expect under the null of no association between transcript abundance and experimental condition. The null distributions for these two kinds of hypothesis tests have often been calculated via permutation, either at the gene-level for the competitive test, or sample

*To whom correspondence should be addressed. ORCID ID: 0000-0003-3174-1656

level, for the self-contained test. Manual permutation is used in an attempt to attain the nominal type 1 error rate^{12,13} or sometimes computationally easier methods that yield similar results^{14,10}. Rejecting a competitive gene-set test is often a higher bar or more difficult burden of proof, and these tests have been more commonly used in gene-set testing. A rejection of a competitive gene-set test can be more biologically meaningful than a self-contained test.

In this paper, we describe within an analytic statistical framework why past gene-set tests have suffered from inflated type 1 error, and how permutation-based methods have sought to address the issue in different ways. Goeman and Buhlmann⁹ made a similar and rigorous description of gene-set tests, though in the context primarily of 2×2 tables, and prior to the development of some of the tests we analyze. More importantly, we additionally describe in the same framework why past and more recently proposed gene-set tests suffer from having variable power as a function of gene set correlation structure.

We show that the location of causal transcripts within the co-expression correlation structure is critical in determining statistical power of the test in addition to the effect sizes of them, and how the issue is relevant to both competitive and self-contained gene-set tests. We define causal transcripts as those expressed genes having a causal effect on the outcome of interest, rather than only being correlated with those that are. Briefly, correlated blocks of genes manifest themselves in disproportionate degree in the tails of distributions under both the null and alternative hypotheses. Since we generally perform statistical tests in the tails of distributions, hypothesis tests are disproportionately influenced by this correlation. When causal transcripts are found in these correlated blocks, power for their detection is likewise overly represented relative to causal transcripts found in less correlated blocks. This observation has been made in the context of region-based SNP testing, though not in the context of gene-set testing¹⁵. Despite the observation being made in the context of SNP-set testing, it is pertinent to mention since tests continue to be proposed that suffer from variable power under different correlation structures¹⁶.

We leverage our statistical analysis of gene-set tests by proposing three tests, the suite of which we call “JAGST” (Just Another Gene Set Test). The first test is self-contained and replicates results from a leading self-contained test ROAST, but does so using a closed form mixture density rather than the rotation and iteration-based method of ROAST. Though ROAST is an elegant, general, and computationally cheap algorithm for various set-based hypothesis tests, two of its more important special cases (mixed and directed set-based tests) fall within our statistical framework so that relevant null tail probabilities can be written down and calculated.

The second two tests of JAGST are power-invariant self-contained and competitive tests. The proposed JAGST power-invariant self-contained test simply uses correlation information of differential expression test statistics to implement a standard chi-square test. The proposed competitive test is more complex and computationally intensive, using a sophisticated and recently proposed high-dimensional penalized regression method for obtaining z-statistics from a variable selection procedure^{17–20}. We use these z-statistics to create an appropriate mixture distribution of non-central chi-squares, which is the null distribution for our competitive test.

We perform simulations with leading and commonly used competitive and self-contained tests, including ROAST, CAMERA, SAFE, and GAGE^{8,14,13,6} to demonstrate the degree to which statistical power is effected by correlation structure of the gene-set and how our proposed tests do not have this same property. ROAST and SAFE are self-contained tests, while CAMERA and GAGE are competitive tests. We perform a data analysis on an already published study of smoking exposure in pregnant women²¹. While expression studies are the context of our analysis, our observations are also relevant to RNAseq. The popular MAST method for RNAseq is based on CAMERA, and many of the same statistical principles will apply in the RNAseq context⁷. Additionally, expression analyses continue to be a valuable and well-used platform for diagnosis and prediction^{22–26}.

2 Methods

Consider matrix Y ($n \times m$) of normalized transcript measures and matrix X ($n \times p$) of experimental conditions, where Y has $m \times m$ correlation matrix Σ , whose i, j^{th} element is $\rho_{i,j}$. Assume for now that $p = 1$ for simplicity. It is common in GSEA to first perform a standard microarray analysis, and then test gene-set enrichment on the summary statistics (such as transcript-level test statistics).

If we fit a regression model of the i^{th} column of Y (Y_i) on X , $i \in 1 \dots m$, we extract some test statistic, call it T_i (oftentimes a t -statistic for linear regression if the outcome is continuous), associated with this model, and likewise T_j for the model of Y_j on X . The test statistic from a linear regression model is

$$T_i \equiv \frac{Y_i^T \cdot X}{\sqrt{\hat{\sigma}_i^2 X^T \cdot X}}$$

if we have centered X and Y , which follows a t -distribution on $n-2$ degrees of freedom under the null hypothesis and where $\hat{\sigma}_i^2$ is the estimated variance of Y_i conditional on X . Likewise, define

$$T_j \equiv \frac{Y_j^T \cdot X}{\sqrt{\hat{\sigma}_j^2 X^T \cdot X}}.$$

Under the null hypothesis of no association between Y_i or Y_j and X , $E[T_i] = E[T_j] = 0$. Since $Cor(\cdot, \cdot)$ is invariant to scaling its arguments and assuming centered, normalized Y_i 's, Y_j 's, and X 's,

$$Cor(T_i, T_j) = E(T_i \cdot T_j) = E(Y_i \cdot X \cdot Y_j \cdot X) = E(Y_i \cdot Y_j) \cdot E(X^2) = \rho_{ij}$$

where the second to last equality holds because we calculate under the null hypothesis of no association with X . So we see that $Cor(T_i, T_j)$ is the (i, j) entry of Σ , the correlation matrix of Y . It follows that the m -length vector T composed of entries T_k , $k = 1, \dots, m$ also has $Cor(T) = \Sigma$.

Now consider the regression model under the alternative hypothesis of an association between an experimental condition and transcript abundance. In particular, suppose there exists a causal association between X and Y_j such that the expectation of the test statistics T_j corresponding to regression model $Y_j \sim X$ is $E[T_j] = \mu_j$. For example, power under this association for an α -level test would be calculated with

$$F_{t_{n-2}}(\mu_j + q) + (1 - F_{t_{n-2}}(\mu_j - q))$$

where $F_{t_{n-2}}(\cdot)$ is the cdf of a t -distribution on $n-2$ degrees of freedom and q is the $\alpha/2$ quantile of that distribution. The expression is the sum of two tail probability regions which corresponds to the rejection region in a two-sided test and is calculated under the alternative hypothesis of $E[T_j] = \mu_j$.

Assume there exists no causal association between X and Y_i , and consider test statistic T_i . Because $Cor(T_i, T_j) = \rho_{i,j}$, $E(T_i) = \rho_{i,j} \cdot \mu_j$; we have power to detect an association between X and Y_i by virtue of Y_i 's correlation with Y_j . though less so because $|\rho_{i,j}| \leq 1$. We will return to this point later in the discussion of statistical power.

Suppose that we have sufficient sample size such that the t -statistics resulting from an expression analysis can be approximated by the standard normal distribution. Since the test statistic for many gene-set tests relies on or is highly correlated with a sum or sum of squares of probe-level summary statistics composing the gene-set (e.g.,^{14,3,8,27,6,13}), consider the following property about the multivariate normal distribution.

$$\sum_{i \in \Gamma} T_i = J_G^T \cdot T \sim N(J_G^T \cdot \mu, J_G^T \Sigma J_G) \quad (1)$$

for T , as above, the length m vector that follows $MVN(\mu, \Sigma)$, Γ a size g set of indices for the gene-set, and J_G a length m indicator vector whose 1 entries correspond to the indices of Γ . So the mean of $\sum_{i \in \Gamma} T_i$ is

the sum of the component means corresponding to the 1 entries in J_G , and the variance of $\sum_{i \in \Gamma} T_i$ is the sum of the entries of Σ that remain after multiplication by J_G . Sometimes this test statistic in the gene-set testing context is called the directed or directional test since its value is sensitive to the sign of T_i .

Additionally, if we square the T_i 's, we have

$$\sum_{i \in \Gamma} T_i^2 = \sum_{i \in \Gamma} \lambda_i \cdot \chi_1^2(\delta_i^2) \text{ where } \delta_i = (\mu_{\mathbf{J}}^T \cdot \nu_i) / \sqrt{\lambda_i}, \quad (2)$$

$\mu_{\mathbf{J}} \equiv J_G \cdot \mu$, and λ_i and ν_i are the eigenvalues and eigenvectors of $J_G^T \Sigma J_G$ ^{28–30}. Sometimes this test statistic in the gene-set testing context is called undirected, non-directional, or mixed as its power is invariant to the sign of T_i .

3 Relationship to permutation-based tests

3.1 Considering preservation of Type 1 error using permutations

It is sometimes thought that permutation is a sure, if computationally burdensome, way of preserving type 1 error when the null distribution of a test statistic is unknown. We show why this is not the case, at least with competitive gene-set tests. The result has been noted frequently elsewhere^{8,27,10}, but we try to describe it in a statistical framework so that evidence of the phenomenon is not primarily simulation-based. First, however, we consider permutation-based self-contained tests and the parametric distributions approximated by them.

3.1.1 Self-contained tests type 1 error

In light of the results in Section 2, we can consider the null distributions generated when samples (“self-contained” gene-set tests) or genes (“competitive” gene-set tests) are permuted to generate a null distribution for the test statistics discussed. In permuting samples, any possible association between gene and outcome is broken so that the expectation of T , the test statistic vector described above, is a length m vector with expectation $\mathbf{0}$, rather than μ . Since the permutations do not permute the particular gene set under consideration, the correlation structure of the gene-set is preserved. Therefore if we define $T_{\{G\}} \equiv J_G^T \cdot T$, one can conclude that a permutation-generated null distribution for $T_{\{G\}}$ is a sample from

$$T_{\{G\},null} \sim MVN(0, \Sigma_{\{G\}}), \text{ where } \Sigma_{\{G\}} \equiv J_G^T \Sigma J_G.$$

So $\Sigma_{\{G\}}$ is the correlation matrix of the transcripts, Y_i 's, in the gene-set. With the joint distribution of $T_{\{G\},null}$ in mind, we can consider the sum of squares of its elements. Since the expectation of the vector $T_{\{G\},null}$ is zero and using notation conventions from Section 2, $\delta_i = 0$ because ν_i will be an inner product with a $\mathbf{0}$ vector. Differences in the null distribution for $\sum_{i \in \Gamma} T_i^2$ are therefore driven by $\{\lambda_i\}$ (the set of eigenvalues of $\Sigma_{\{G\}}$), and the null distribution can be very different depending on the distribution of the $\{\lambda_i\}$'s. If there is high within-set correlation, greater variation in λ_i will lead to a heavier-tailed null distribution, whereas little within-set correlation will correspond to relatively similar λ_i 's and a thinner tail.

No within gene-set correlation implies independent T_i 's, so that

$$\sum_{i=1}^n T_i^2 \sim \chi_g^2.$$

It is because independent component test statistics have been assumed at times in the past that a χ_g^2 null distribution is used when in fact the heavier tailed $\sum_{i=1}^n \lambda_i \cdot X_1^2$ was the true null distribution⁴. It is in this distinct way that even self-contained tests have occasionally had inflated type 1 error, though this has been recognized and corrected in different articles^{31,14,10,32}.

3.1.2 Competitive tests type 1 error

When permuting across gene sets for a competitive gene set test, probe-wise associations with the outcome are maintained, but the set over which we aggregate changes as a function of the genes randomly chosen to compose our permutation-generated set. Thus, for permutation k of the null distribution generation, the distribution from which we sample is

$$T_{k,null,comp} \sim MVN(\mu_{\{G_k\}}, \Sigma_{\{G_k\}}),$$

where $\{G_k\}$ is a set of genes of size g randomly chosen, $\mu_{\{G_k\}}$ is the expectation vector of test statistics corresponding to those particular genes, and $\Sigma_{\{G_k\}}$ is a submatrix of Σ corresponding to the gene set $\{G_k\}$. Since $\{G_k\}$ changes with each iteration, it is evident that $\{T_{1,null,comp}, \dots, T_{q,null,comp}\}$, where q is the size of the permutation-generated null, are not identically distributed, but a mixture of distributions.

Crucially, and likely the reason type 1 error has been difficult to maintain in competitive gene-set tests, the off-diagonal entries of $\Sigma_{\{G_k\}}$ (the correlation matrix of permutation k) will often be systematically smaller than the off-diagonal entries of $\Sigma_{\{G_k\}}$ because genes within biologically meaningful sets will often be co-expressed to a greater degree than randomly chosen groups of genes. We may therefore expect that the permutation-generated null distribution will not be as thick tailed as the test statistic whose null it is trying to approximate.

3.2 Statistical power when using permutations

Because expression analyses often rely on univariate regression test statistics, confounding structures are not detected and power to detect some non-causal transcripts can be above α because of their association with the causal gene. The underlying causal structure of transcript on phenotype and confounding by co-expressed genes and their effects on statistical power has been less discussed on a methodological level. When variable power has been discussed, it has been done so from the perspective of underlying biological or platform phenomena, such as in Oshlack and Wakefield³³ with respect to transcript length bias. Here we describe why statistical power for both self-contained and competitive gene-set tests can vary due to the methodological set-up of the tests themselves.

3.2.1 Illustrative example

We first consider an example to explain variable power more concretely. For simplicity, assume a gene-set consists of 3 genes, one of which has a causal association with the outcome, and the other two do not. Call the test statistic associated with the causal gene T_1 with expectation μ_1 , and those of the non-causal genes T_2 and T_3 (both with expectations 0). First suppose that T_2 and T_3 have correlation ρ , but both are independent of T_1 (the causal gene). The distribution of their sum is

$$T_1 + T_2 + T_3 \sim N(\mu_1, 3 + 2\rho),$$

while the non-directional test statistic is

$$T_1^2 + T_2^2 + T_3^2 \sim \sum_{i=1}^3 \lambda_i \cdot \chi_1^2(\delta_i^2) \text{ with } \delta_i = (\mu_1 \ 0 \ 0) \cdot \nu_i$$

and λ_i and ν_i are the eigenvalues and eigenvectors of Σ .

Consider now the scenario where we “shift” the correlation structure so that T_1 and T_2 have correlation ρ , and T_3 is independent of them both. though the same causal associations remain. So $E(T_1) = \mu_1$, and in

this scenario $E(T_2) = \mu_1\rho$ because of T_2 's correlation with T_1 . The expectation of T_3 is still 0. This time, the distribution of their sum is

$$T_1 + T_2 + T_3 \sim N(\mu_1(1 + \rho), 3 + 2\rho)$$

and the non-directional test statistic is

$$T_1^2 + T_2^2 + T_3^2 \sim \sum_{i=1}^3 \lambda_i \cdot \chi_1^2(\delta_i^{*2}) \text{ where } \delta_i^* = (\mu_1 \quad \mu_1\rho \quad 0) \cdot \nu_i.$$

It will always be the case that $\delta_i^{*2} \geq \delta_i^2 \forall i$, that is, in the second scenario the non-centrality parameters will tend to be bigger than those in the first scenario. We know this because $(\mu_1 \quad \mu_1\rho \quad 0)$ is greater than $(\mu_1 \quad 0 \quad 0)$ element-wise. Additionally, there will exist at least one δ_i strictly greater in the second scenario than the first scenario because for at least one i , ν_i has a non-zero entry in the second element since eigenvectors span the space defined by the columns of Σ . Because χ^2 distributions are stochastically strictly increasing in their non-centrality parameters, we will always be better powered to detect the gene-set in the second scenario than in the first.

From this example, we conclude that we are more likely to reject the null for gene set enrichment test when causal transcripts are co-expressed with non-causal ones.

4 Proposed hypothesis tests: JAGST

4.1 Proposed hypothesis test 1: an analytic approach to ROAST's self-contained test

First we introduce directed and mixed versions of self-contained gene-set test that are nearly numerically equivalent to these important special cases of ROAST's versions of them¹⁴. Our directed and mixed tests rely on closed form formulae rather than the rotation-based methodology of ROAST, however. ROAST is itself an elegant way of approximating the results of permutation-based procedures.

The test statistics for the directed and mixed self-contained gene-set tests were in fact already introduced in Section 2 with equations (1) and (2), respectively. Equation (1) is the sum of differential expression test statistics, while equation (2) is the sum of squares of differential expression test statistics. Since the null distributions of (1) and (2) are known and provided in the expressions, we do not need to rely on the rotation of residual methodology, which though statistically elegant has p-value granularity dependent on the number of specified iterations. We can instead calculate p-values using tail probabilities of the normal distribution in the case of the directional test (eqn (1)) and of a mixture of scaled χ^2 distributions in the case of the mixed test (eqn (2)).

We see in the Supplementary Material the correlation of 0.99 (0.96) between the ROAST directional test (mixed test) and analytic calculation of tail probabilities using these densities, both under the null hypothesis. That correlation falls slightly under the alternative hypothesis for reasons explained in the figures, but remains above 0.88. If the variable power issue is not a concern for the analyst, for example for reasons given in the Discussion section, and p-values from the self-contained test ROAST methodology are preferred, one need use only the test statistics and explicit null distributions in (1) or (2). Since inversion of the χ^2 mixture distribution can be difficult, sampling from the mixture distribution is also adequate and likely still more computationally inexpensive than ROAST.

4.2 Proposed hypothesis test 2: power-invariant competitive gene-set test

We propose the JAGST competitive test whose statistical power is invariant to the correlation structure in which causal effects are found. We do so by making two key changes to currently used gene-set testing methods. First, define $T_{\{G\}}$ and $\Sigma_{\{G\}}$ as the vector of differential expression test statistics corresponding

to gene set $\{G\}$, and the correlation matrix of the Y_i 's making up the gene-set, respectively. Then the test statistic for the competitive test is calculated using

$$T_{\{G\}}^T \Sigma_{\{G\}}^{-1} T_{\{G\}},$$

rather than a straightforward sum or sum of squares of test statistics. Using the inverse covariance matrix of the test statistics is central to achieving the desired correlation structure power invariance. Secondly, the null distribution is calculated by taking random subsets of genes of size g and calculating the same test statistic. Unlike other competitive tests that calculate their null distribution with random subsets of genes, using the inverse correlation matrix makes each realization of our null distribution a function of only the underlying effects of each subset thereby controlling for correlation structure.

While this is the essential idea of the JAGST competitive test, in practice we generate the proposed null distribution in a different and computationally easier way. Especially for large gene-sets, some of which approach a few hundred genes, calculating the null distribution in this brute force way is too cumbersome if not numerically prohibitive depending on whether the correlation matrices are approximately singular.

We therefore propose instead to:

1. Perform variable selection on the random gene subsets of size g ,
2. choose the model with the smallest AIC (or another model fitness criterion that could be shown to control type 1 error and be power-invariant), then
3. use a novel and sophisticated method to calculate the z-statistics of the selected variables^{19,18,20}, in order to
4. take the sum of squares of these test statistics, calling them $\delta_{PI} \equiv \sum \delta_{i,PI}^2$ ("PI" for "power-invariant"), and
5. sample from $\chi_g^2(\delta_{PI})$, which is the null distribution of the above test statistic as demonstrated in simulation (Figure 4).

We iterate until desired granularity in the null distribution is achieved.

The L1-penalized regression hypothesis testing procedure of Taylor and Tibshirani¹⁹ proves central to our algorithm. It provides a means of obtaining z-statistics on selected variables in a penalized regression framework and until recently was not possible. These z-statistics allow us to determine the distribution under the alternative hypothesis that our differential expression test statistics arise from.

While variable selection is not without computation cost, since the procedure stops once a certain AIC is achieved, we will generally avoid the large $O(g^3)$ cost of matrix inversion and numeric instability for large gene sets by a significant margin. This is especially true if it is assumed that in general gene sets are composed of many transcripts working together in pathways, though reveal relatively sparse models when regularization is applied.

Iteration over different sets of size g will yield a mixture of non-central χ^2 distributions. Formulae and approximations have been proposed for mixtures of χ^2 distributions^{29,34}, which could be used if the same null distribution were relevant to many different tests, or sampling from distributions were burdensome. For our purposes, we eschew these formulae to prefer sampling from the mixture distribution in R³⁵.

4.3 Proposed hypothesis test 3: power-invariant self-contained gene-set test

For the JAGST self-contained test, we again use the test statistic

$$T_{\{G\}}^T \Sigma_{\{G\}}^{-1} T_{\{G\}},$$

with notation used as before. In this case however, the null hypothesis assumes no association between transcript and outcome so the null distribution for this test statistic is simply χ_g^2 . The quantiles for this null distribution will often be a much lower threshold for statistical significance as compared to the proposed competitive gene set test.

5 Results

We provide simulation and data analysis results for the different JAGST tests and compare them with CAMERA (with and without ranks), GAGE (with and without ranks), ROAST, and SAFE. We abbreviate our simulation analysis of type 1 error since it has been described well and studied in detail elsewhere^{13,36,10,8} and focus on power as a function of correlation structure under different generating models. We perform simulation and data analysis in the R language³⁵.

5.1 Simulation

We include a detailed description of our simulation analysis in the *Supplementary Material* of the manuscript. Here, we briefly describe that for each of ROAST, GAGE, SAFE, and CAMERA, we produce power curves for when causal transcripts are found in areas of correlation (dominating curve in Figures 1-6) or lack of correlation (lower curve in Figures 1-6). Under the same simulation scenario, we applied the JAGST competitive (Figure 7) and self-contained tests (Figure 8), and noted no power curve divergence. This indicates that there is the same power for the statistical test regardless of whether transcripts are found in correlated regions.

We also performed gene-set tests under the null hypothesis (Figure 9) and weak alternative (Figure 10) to compare results from ROAST's directional tests with tail probabilities justified in Section 2. Comparison of resultant p-values are plotted against one another in these figures. Additional description of this simulation is given in the *Supplementary Material*.

5.2 Data analysis

We analyze data from an already published study of maternal and fetal transcription variation due to smoking exposure²¹. The study analyzed maternal peripheral, placenta, and neonatal cord blood on 20 pregnant women with smoking exposure and 50 without significant smoking exposure. Women were assayed using Illumina expression beadchip v3 and yielded 24,526 transcripts. We accessed the data via its Gene Expression Omnibus (GEO) accession number GDS3929^{37,38}.

In keeping with the analysis protocol of Votavova et al.²¹, we filtered transcripts of each cell type to approximately the third with a detectable level of transcript abundance. In doing so, the differential expression analysis on the remaining genes was consistent with that of²¹ thus confirming data processing and integrity. For our gene-set analyses, we used only this approximate third most differentially expressed in each cell type category.

We analyzed 7 different gene-sets, motivated by ones commonly used in gene-set enrichment studies, oncogenic pathways, and ones more specific to the context of^{39-41,21} (e.g., placenta development). We used 8 different gene-set test methodologies: CAMERA (with and without ranks), ROAST, SAFE, and GAGE (with and without ranks). ROAST and SAFE are self-contained tests, while CAMERA and GAGE are competitive tests. The last two used are our proposed power-invariant competitive and self-contained tests. P-values for the 8 tests by gene-set and data source (maternal peripheral blood, neonatal cord blood, or placenta) are included in Tables 1, 2, and 3, respectively. The JAGST competitive and self-contained tests are some of the least significant tests in the cases of nearly every gene-set and cell type category. This suggests that causal transcripts are correlated with many other genes in the gene-set, and our proposed tests successfully control for the correlation whereas the other gene-set tests are affected by it.

6 Discussion

Much has been learned from gene-set testing since the early 2000's when such methods were developed. The idea that because biological processes are complex systems whose elements are not acting in isolation is powerful and has rightly been leveraged in many gene-set testing methodologies. Most of these methods are uniquely suited to certain situations, implicitly or explicitly making different assumptions of the underlying data and biological process.

We described many gene-set tests within an explicit statistical framework in this paper and showed why some could be subject to type 1 error inflation. While that observation has been made previously, its best description was done in the context of 2 x 2 tables⁹, and we describe why it can be the case with a linear combination of differential expression statistics.

More importantly in this paper, we observed and explained why all gene-set tests analyzed have variable power as a function of correlation structure. We show that if gene drivers of phenotype are found in uncorrelated regions of a gene-set, there is much less power to detect the gene-set than if the drivers are in a correlated region of the gene-set. The observation is particularly relevant because co-expression of genes on expression arrays is ubiquitous and variable; indeed, the very idea of gene-set testing is that groups of genes working in concert are expressed and can be tested together.

The gene set tests applied in our data analysis of smoking exposure on pregnant women did not yield any significant results adjusted for multiple testing. Our proposed self-contained and competitive tests were less significant in the cases of nearly every gene set, which is consistent with the idea that causal transcripts in the gene-sets tested are often correlated with non-causal ones. Since our tests control for that correlation, there is no more power for detection than if the causal transcripts were independent of non-causal ones. The proposed JAGST competitive gene set test is less significant than the JAGST self-contained test, which is consistent with the idea that competitive tests are a more significant burden of proof than self-contained tests. While we were able to replicate the differential expression results of²¹, we were not able to confirm findings of significant gene-sets. Since²¹ relied on the DAVID database for enrichment analyses, we cannot expect our results to coincide with those of the previous study^{42,43}.

The self-contained tests used in our data analysis, ROAST and SAFE, were not generally more significant than the competitive tests, Camera and Gage. Since tests are a function of correlation structure, and in the cases of the ranked tests, non-parametric, we would not necessarily expect an obvious distinction between the two kinds of tests even though their null hypotheses are different.

It is unclear why exactly there is a small divergence in power curves for the JAGST competitive test in the simulation scenarios compared (see Figure 7). Since our method relies on test statistics from an L1-penalized regression model, it is likely the relatively small variation is due to very different correlation structures from the simulation in which the transcripts with non-zero effect sizes are found. These correlation structures may affect AIC, the model fitness criterion being used to choose the tuning parameter, which then affects magnitude of the test statistics. It is important, though, that that small difference we see in our test is much smaller than that observed in other tests and that additionally the two power curves are parallel on the -log p-value scale. In contrast, we see that divergence in power curves increases with effect size for the other tests examined. For the JAGST self-contained test, we observe nearly identical power curves indicating nearly perfect power invariance.

7 Software and Data

An R package for implementing JAGST is available on GitHub under “sojourningNorth/JAGST” and can be installed in the standard way using directions found at the repository. Additionally, code used for simulation and analysis found in the manuscript is available at “sojourningNorth/JAGST-analysis”. Data used for analysis is available through NCBI's GEO database and has accession number GDS3929^{37,38}.

8 Supplementary Material

A detailed description of our simulation analysis as well as figures showing correspondences between the directional and mixed tests of ROAST and our analytical calculation of them can be found in the *Supplementary Material*.

Acknowledgements

The author wishes to thank Prof. Arnaldo Frigessi for helpful comments in the preparation of the manuscript.

References

- [1] Subramanian, A., Tamayo, P., Mootha, V. K., Mukherjee, S., Ebert, B. L., Gillette, M. A., Paulovich, A., Pomeroy, S. L., Golub, T. R., Lander, E. S., and Mesirov, J. P. Gene set enrichment analysis: A knowledge-based approach for interpreting genome-wide expression profiles. *Proceedings of the National Academy of Sciences*, 102(43):15545–15550, 2005. doi: 10.1073/pnas.0506580102. URL <http://www.pnas.org/content/102/43/15545.abstract>.
- [2] Goeman, J. J., Van De Geer, S. A., De Kort, F., and Van Houwelingen, H. C. A global test for groups of genes: testing association with a clinical outcome. *Bioinformatics*, 20(1):93–99, 2004.
- [3] Kim, S.-Y. and Volsky, D. J. PAGE: parametric analysis of gene set enrichment. *BMC bioinformatics*, 6:144, 2005. ISSN 1471-2105. doi: 10.1186/1471-2105-6-144.
- [4] Irizarry, R. a., Wang, C., Zhou, Y., and Speed, T. P. Gene Set Enrichment Analysis Made Simple. *Statistical methods for medical research*, 18(6):565–575, 2009. ISSN 1477-0334. doi: 10.1177/0962280209351908.Gene.
- [5] Efron, B. and Tibshirani, R. On testing the significance of sets of genes. *The Annals of Applied Statistics*, 1(1):107–129, 2007. ISSN 1932-6157. doi: 10.1214/07-AOAS101. URL <http://projecteuclid.org/euclid.aoas/1183143731>.
- [6] Luo, W., Friedman, M. S., Shedden, K., Hankenson, K. D., and Woolf, P. J. GAGE: generally applicable gene set enrichment for pathway analysis. *BMC bioinformatics*, 10:161, 2009. ISSN 1471-2105. doi: 10.1186/1471-2105-10-161.
- [7] Finak, G., McDavid, A., Yajima, M., Deng, J., Gersuk, V., Shalek, A. K., Slichter, C. K., Miller, H. W., McElrath, M. J., Prlic, M., Linsley, P. S., and Gottardo, R. MAST: a flexible statistical framework for assessing transcriptional changes and characterizing heterogeneity in single-cell RNA sequencing data. *Genome Biology*, 16(1):278, 2015. ISSN 1474-760X. doi: 10.1186/s13059-015-0844-5.
- [8] Wu, D. and Smyth, G. K. Camera: A competitive gene set test accounting for inter-gene correlation. *Nucleic Acids Research*, 40(17):1–12, 2012. ISSN 03051048. doi: 10.1093/nar/gks461.
- [9] Goeman, J. J. and Buhlmann, P. Analyzing gene expression data in terms of gene sets: Methodological issues. *Bioinformatics*, 23(8):980–987, 2007. ISSN 13674803. doi: 10.1093/bioinformatics/btm051.
- [10] Zhou, Y. H., Barry, W. T., and Wright, F. A. Empirical pathway analysis, without permutation. *Biostatistics*, 14(3):573–585, 2013. ISSN 14654644. doi: 10.1093/biostatistics/kxt004.
- [11] Rahmatallah, Y., Emmert-Streib, F., and Glazko, G. Gene set analysis for self-contained tests: Complex null and specific alternative hypotheses. *Bioinformatics*, 28(23):3073–3080, 2012. ISSN 13674803. doi: 10.1093/bioinformatics/bts579.

- [12] Naeem, H., Zimmer, R., Tavakkolkhah, P., and Küffner, R. Rigorous assessment of gene set enrichment tests. *Bioinformatics*, 28(11):1480–1486, 2012. ISSN 13674803. doi: 10.1093/bioinformatics/bts164.
- [13] Barry, W. T., Nobel, A. B., and Wright, F. A. Significance analysis of functional categories in gene expression studies: A structured permutation approach. *Bioinformatics*, 21(9):1943–1949, 2005. ISSN 13674803. doi: 10.1093/bioinformatics/bti260.
- [14] Wu, D., Lim, E., Vaillant, F., Asselin-Labat, M. L., Visvader, J. E., and Smyth, G. K. ROAST: Rotation gene set tests for complex microarray experiments. *Bioinformatics*, 26(17):2176–2182, 2010. ISSN 13674803. doi: 10.1093/bioinformatics/btq401.
- [15] Swanson, D., Blacker, D., AlChawa, T., Ludwig, K., Mangold, E., and Lange, C. Properties of permutation-based gene tests and controlling type 1 error using a summary statistic based gene test. *BMC Genetics*, 14(1):108, 2013. ISSN 1471-2156. doi: 10.1186/1471-2156-14-108.
- [16] Bakshi, A., Zhu, Z., Vinkhuyzen, A. A., Hill, W. D., McRae, A. F., Visscher, P. M., and Yang, J. Fast set-based association analysis using summary data from gwas identifies novel gene loci for human complex traits. *Scientific reports*, 6:32894, 2016.
- [17] Zou, H., Hastie, T., and Tibshirani, R. On the "degrees of freedom" of the lasso. *Annals of Statistics*, 35(5):2173–2192, 2007. ISSN 00905364. doi: 10.1214/009053607000000127.
- [18] Lockhart, R., Taylor, J., Tibshirani, R. J., and Tibshirani, R. A significance test for the lasso. *Annals of Statistics*, 42(2):413–468, 2014. ISSN 00905364. doi: 10.1214/13-AOS1175.
- [19] Taylor, J. and Tibshirani, R. Post-selection inference for ℓ_1 -penalized likelihood models. *Canadian Journal of Statistics*, (2015):1–26, 2017. ISSN 1708945X. doi: 10.1002/cjs.11313.
- [20] Tibshirani, R. J., Taylor, J., Lockhart, R., and Tibshirani, R. Exact Post-Selection Inference for Sequential Regression Procedures. *Journal of the American Statistical Association*, 111(514):600–620, 2016. ISSN 0162-1459. doi: 10.1080/01621459.2015.1108848.
- [21] Votavova, H., Dostalova Merkerova, M., Fejglova, K., Vasikova, A., Krejcik, Z., Pastorkova, A., Tabashidze, N., Topinka, J., Veleminsky, M., Sram, R. J., and Brdicka, R. Transcriptome alterations in maternal and fetal cells induced by tobacco smoke. *Placenta*, 32(10):763–770, 2011. ISSN 01434004. doi: 10.1016/j.placenta.2011.06.022. URL <http://dx.doi.org/10.1016/j.placenta.2011.06.022>.
- [22] Tachibana, C. Transcriptomics today: Microarrays, RNA-seq, and more. *Science*, 2015. ISSN 0036-8075. doi: 10.1126/science.opms.p1500095.
- [23] Zhao, S., Fung-Leung, W. P., Bittner, A., Ngo, K., and Liu, X. Comparison of RNA-Seq and microarray in transcriptome profiling of activated T cells. *PLoS ONE*, 9(1), 2014. ISSN 19326203. doi: 10.1371/journal.pone.0078644.
- [24] Byron, S. A., Van Keuren-Jensen, K. R., Engelthaler, D. M., Carpten, J. D., and Craig, D. W. Translating RNA sequencing into clinical diagnostics: opportunities and challenges. *Nature Reviews Genetics*, 17(5):257–271, 2016. ISSN 1471-0056. doi: 10.1038/nrg.2016.10.
- [25] Zhang, W., Yu, Y., Hertwig, F., Thierry-Mieg, J., Zhang, W., Thierry-Mieg, D., Wang, J., Furlanello, C., Devanarayan, V., Cheng, J., Deng, Y., Hero, B., Hong, H., Jia, M., Li, L., Lin, S. M., Nikolsky, Y., Oberthuer, A., Qing, T., Su, Z., Volland, R., Wang, C., Wang, M. D., Ai, J., Albanese, D., Asgharzadeh, S., Avigad, S., Bao, W., Bessarabova, M., Brilliant, M. H., Brors, B., Chierici, M., Chu, T.-M., Zhang, J., Grundy, R. G., He, M. M., Hebbring, S., Kaufman, H. L., Lababidi, S., Lancashire, L. J., Li, Y., Lu, X. X., Luo, H., Ma, X., Ning, B., Noguera, R., Peifer, M., Phan, J. H., Roels, F., Rosswog, C., Shao, S., Shen, J., Theissen, J., Tonini, G. P., Vandesompele, J., Wu, P.-Y., Xiao, W.,

- Xu, J., Xu, W., Xuan, J., Yang, Y., Ye, Z., Dong, Z., Zhang, K. K., Yin, Y., Zhao, C., Zheng, Y., Wolfinger, R. D., Shi, T., Malkas, L. H., Berthold, F., Wang, J., Tong, W., Shi, L., Peng, Z., and Fischer, M. Comparison of RNA-seq and microarray-based models for clinical endpoint prediction. *Genome biology*, 16:133, 2015. ISSN 1474-760X. doi: 10.1186/s13059-015-0694-1.
- [26] Xu, J., Gong, B., Wu, L., Thakkar, S., Hong, H., and Tong, W. Comprehensive assessments of RNA-seq by the SEQC consortium: FDA-led efforts advance precision medicine. *Pharmaceutics*, 8(1), 2016. ISSN 19994923. doi: 10.3390/pharmaceutics8010008.
- [27] Ritchie, M. E., Phipson, B., Wu, D., Hu, Y., Law, C. W., Shi, W., and Smyth, G. K. limma powers differential expression analyses for RNA-sequencing and microarray studies. *Nucleic acids research*, 43(7):e47, 2015. ISSN 13624962. doi: 10.1093/nar/gkv007.
- [28] Imhof, J. P. Computing the distribution of quadratic forms in normal variables. *Biometrika*, 48(3/4): 419–426, 1961.
- [29] Press, S. J. Linear combinations of non-central chi-square variates. *The Annals of Mathematical Statistics*, pages 480–487, 1966.
- [30] Harville, D. On the Distribution of Linear Combinations of Non-central Chi-Squares. *The Annals of Mathematical Statistics*, 42(2):809–811, 1971.
- [31] Gatti, D. M., Barry, W. T., Nobel, A. B., Rusyn, I., and Wright, F. A. Heading down the wrong pathway: on the influence of correlation within gene sets. *BMC genomics*, 11(1):574, 2010. ISSN 1471-2164. doi: 10.1186/1471-2164-11-574.
- [32] Mesirov, P. T., Steinhardt, G., Liberzon, A., and P, J. The limitations of simple gene set enrichment analysis assuming gene independence. *Statistical Methods in Medical Research*, 25(1):472–487, 2016.
- [33] Oshlack, A. and Wakefield, M. Transcript length bias in RNA-seq data confounds systems biology. *Biology direct*, 4:14, 2009. ISSN 1745-6150. doi: 10.1186/1745-6150-4-14. URL <http://www.ncbi.nlm.nih.gov/pubmed/19371405>.
- [34] Liu, H., Tang, Y., and Zhang, H. H. A new chi-square approximation to the distribution of non-negative definite quadratic forms in non-central normal variables. *Computational Statistics and Data Analysis*, 53(4):853–856, 2009. ISSN 01679473. doi: 10.1016/j.csda.2008.11.025. URL <http://dx.doi.org/10.1016/j.csda.2008.11.025>.
- [35] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2017. URL <https://www.r-project.org/>.
- [36] Barry, W. T., Nobel, A. B., and Wright, F. A. A statistical framework for testing functional categories in microarray data. *Annals of Applied Statistics*, 2(1):286–315, 2008. ISSN 19326157. doi: 10.1214/07-AOAS146.
- [37] Edgar, R., Domrachev, M., and Lash, A. E. Gene Expression Omnibus: NCBI gene expression and hybridization array data repository. *Nucleic acids research*, 30(1):207–210, 2002.
- [38] Barrett, T., Wilhite, S. E., Ledoux, P., Evangelista, C., Kim, I. F., Tomashevsky, M., Marshall, K. A., Phillippy, K. H., Sherman, P. M., Holko, M., Yefanov, A., Lee, H., Zhang, N., Robertson, C. L., Serova, N., Davis, S., and Soboleva, A. NCBI GEO: Archive for functional genomics data sets - Update. *Nucleic Acids Research*, 41(D1):991–995, 2013. ISSN 03051048. doi: 10.1093/nar/gks1193.
- [39] Nevins, J. R. and Potti, A. Mining gene expression profiles: expression signatures as cancer phenotypes. *Nature reviews. Genetics*, 8(8):601–609, 2007. ISSN 1471-0056. doi: 10.1038/nrg2137. URL wos:000248170700014.

- [40] Zhang, J., Chen, Y.-H., and Lu, Q. Pro-oncogenic and anti-oncogenic pathways: opportunities and challenges of cancer therapy. *Future oncology (London, England)*, 6(4):587–603, 2010. ISSN 1479-6694. doi: 10.2217/fon.10.15.
- [41] Bild, A. H., Yao, G., Chang, J. T., Wang, Q., Potti, A., Chasse, D., Joshi, M. B., Harpole, D., Lancaster, J. M., Berchuck, A., Olson Jr., J. A., Marks, J. R., Dressman, H. K., West, M., and Nevins, J. R. Oncogenic pathway signatures in human cancers as a guide to targeted therapies. *Nature*, 439(7074):353–357, 2006. ISSN 0028-0836. doi: 10.1038/nature04296. URL <http://www.ncbi.nlm.nih.gov/pubmed/16273092>.
- [42] Huang, D. W., Sherman, B. T., and Lempicki, R. A. Systematic and integrative analysis of large gene lists using DAVID bioinformatics resources. *Nature Protocols*, 4(1):44–57, 2008. ISSN 1754-2189. doi: 10.1038/nprot.2008.211. URL <http://www.nature.com/doifinder/10.1038/nprot.2008.211>.
- [43] Huang, D. W., Sherman, B. T., and Lempicki, R. A. Bioinformatics enrichment tools: Paths toward the comprehensive functional analysis of large gene lists. *Nucleic Acids Research*, 37(1):1–13, 2009. ISSN 03051048. doi: 10.1093/nar/gkn923.

Table 1: P-values from different gene-set tests on maternal peripheral blood data.

Pathway	p53	MAPK	TGF_beta	cell_prog	plac	inflamm	immune
Camera (ranks)	0.43	0.22	0.98	0.46	0.25	0.48	0.35
Camera	0.64	0.26	0.91	0.50	0.22	0.57	0.32
Roast	0.95	0.17	0.66	0.90	0.43	0.35	0.92
Safe	0.25	0.07	0.33	0.62	0.37	0.04	0.23
Gage (ranks)	0.14	0.13	0.34	0.36	0.28	0.06	0.12
Gage	0.08	0.13	0.22	0.32	0.42	0.07	0.12
JAGST comp. test	0.65	1.00	0.68	1.00	0.88	1.00	1.00
JAGST self-cont. test	0.21	1.00	0.26	0.97	0.58	1.00	1.00

Table 2: P-values from different gene-set tests on neonatal cord blood data.

Pathway	p53	MAPK	TGF_beta	cell_prog	plac	inflamm	immune
Camera (ranks)	0.76	0.69	0.61	0.83	0.38	0.54	0.67
Camera	0.93	0.71	0.66	0.88	0.45	0.59	0.70
Roast	0.84	0.93	0.63	0.77	0.66	0.81	0.91
Safe	0.47	0.55	0.65	0.85	0.31	0.12	0.75
Gage (ranks)	0.46	0.58	0.50	0.92	0.49	0.17	0.79
Gage	0.60	0.57	0.40	0.85	0.39	0.11	0.77
JAGST comp. test	0.96	1.00	0.58	1.00	0.27	1.00	1.00
JAGST self-cont. test	0.85	1.00	0.24	0.99	0.07	1.00	1.00

Table 3: P-values from different gene-set tests on placenta data.

Pathway	p53	MAPK	TGF_beta	cell_prog	plac	inflamm	immune
Camera (ranks)	0.10	0.86	0.45	0.27	0.09	0.34	0.98
Camera	0.10	0.95	0.49	0.21	0.07	0.20	0.91
Roast	0.23	0.33	0.78	0.41	0.18	0.04	0.43
Safe	0.18	0.03	0.15	0.35	0.75	0.11	0.31
Gage (ranks)	0.23	0.05	0.21	0.52	0.81	0.12	0.45
Gage	0.26	0.04	0.23	0.49	0.77	0.07	0.41
JAGST comp. test	0.89	1.00	0.89	1.00	0.76	1.00	1.00
JAGST self-cont. test	0.66	1.00	0.65	1.00	0.43	1.00	1.00

Figure 1: Comparison of power between SAFE tests in correlated region (solid line) and uncorrelated (long dashes) region as a function of the strength of association between two transcripts within the gene set and the outcome, holding the sample size constant at 280 observations. There is a practical ceiling on the solid line curve because an empirical p-value is calculated and iterations were limited to 1000.

Figure 2: Comparison of power between ROAST tests in correlated region (solid line) and uncorrelated (long dashes) region as a function of the strength of association between two transcripts within the gene set and the outcome, holding the sample size constant at 280 observations. There is a practical ceiling on the solid line curve because an empirical p-value is calculated and iterations were limited to 1000.

Figure 3: Comparison of power between CAMERA tests in correlated region (solid line) and uncorrelated (long dashes) region as a function of the strength of association between two transcripts within the gene set and the outcome, holding the sample size constant at 280 observations.

Figure 4: Comparison of power between CAMERA tests using ranks in correlated region (solid line) and uncorrelated (long dashes) region as a function of the strength of association between two transcripts within the gene set and the outcome, holding the sample size constant at 280 observations.

Figure 5: Comparison of power between GAGE tests in correlated region (solid line) and uncorrelated (long dashes) region as a function of the strength of association between two transcripts within the gene set and the outcome, holding the sample size constant at 280 observations.

Figure 6: Comparison of power between GAGE tests using ranks in correlated region (solid line) and uncorrelated (long dashes) region as a function of the strength of association between two transcripts within the gene set and the outcome, holding the sample size constant at 280 observations.

Figure 7: Comparison of power between the JAGST competitive test in correlated region (solid line) and uncorrelated (long dashes) region as a function of the strength of association between two transcripts within the gene set and the outcome, holding the sample size constant at 280 observations.

Figure 8: Comparison of power between the JAGST self-contained test in correlated region (solid line) and uncorrelated (long dashes) region as a function of the strength of association between one transcript within the gene set and the outcome, holding the sample size constant at 280 observations.

Figure 9: Comparison of p-values between ROAST's directional test and tail probabilities calculated analytically. These p-values were generated under the null and have a correlation of 0.99. The line corresponds to $y=x$, along which equivalent points will fall.

Figure 10: Comparison of p-values between ROAST's directional test and tail probabilities calculated analytically. These p-values were generated under a weak alternative and have a correlation of 0.99. The line corresponds to $y=x$, along which equivalent points will fall.

Statistical power of gene-set enrichment analysis is a function of gene set correlation structure: Supplementary Material

David Swanson

S2 Simulation

We generated 280 samples, each with 40 transcripts or features, the 40 composed of a correlated region of size 20 with correlation 0.8 and an uncorrelated region of size 20. The correlated and uncorrelated blocks also compose our two hypothetical gene sets. We then generated two different binary outcomes consistent with a logistic regression model, where the probability of event in one case was a function of two transcripts in the correlated block and in the other case a function of two uncorrelated transcripts in the uncorrelated block. The effect sizes of the two transcripts were equal and were also the same in either scenario (i.e., whether the causal transcripts were in the correlated or uncorrelated blocks).

We varied the effect size of the two transcripts from a log odds ratio of zero to 3 in the correlated region and performed eight different gene set tests on the 20 features composing the correlated block: CAMERA (with and without ranks), GAGE (with and without ranks), ROAST, SAFE, and the two JAGST power-invariant tests were proposed, one self-contained and the other competitive.

We then varied the effect sizes of the two transcripts in the uncorrelated block over the same values and performed the same eight gene tests, this time on the 20 features composing the uncorrelated block. The two power curves in each of Figures 1-6 found in the main manuscript correspond to the respective test applied to the 20 correlated or uncorrelated features and the associated outcome.

In the cases of CAMERA, ROAST, SAFE, and GAGE, at all effect sizes greater than zero, there was more power to detect the correlated block gene set, even though effect sizes were the same as compared to the uncorrelated block gene set (Figures 1-6). Indeed, the power curves diverged and had different slopes on the $-\log$ p-value scale so that there was a still bigger power difference at increasing effect sizes. The divergence is particularly striking with GAGE, whose simulation was repeated to confirm the result. In the cases of Figures 1 and 2, we see a ceiling on the $-\log$ p-value, essentially stopping the divergence of power curves. This occurs because the methods use empirical p-values, and we set the number of iterations to 1000. The divergence would continue if the number of iterations were increased. There is a similar ceiling on the rank-based tests found in Figures 4 and 6.

Our proposed gene set tests on the other hand yielded power curves that were much closer at all effect sizes of the two underlying causal transcripts (Figures 7 and 8). While there are small power differences with at least the proposed competitive JAGST test, the power curves were parallel, indicating that the power difference in curves does not increase for larger effect sizes.

Lastly we performed gene-set tests under the null hypothesis and weak alternative to compare results from ROAST's directional tests with tail probabilities justified in Section 2 of the main manuscript. The alternative hypothesis posited two causal transcripts among a gene-set size of 20 correlated transcripts. Rather than inverting the distribution, we estimated this tail probability using the appropriate mixture distribution of central, scaled χ^2 's. Figures 9 and 10 show the high degree of correspondence between the two ways of calculating the test, with a correlation of 0.99. Figures found in the Supplementary Material show the correspondence between the mixed, non-directional tests of ROAST and the analytic solution described in Section 2. Again there is a high degree of correspondence between the two tests under the null and alternative hypotheses, both with correlations above 0.9.

S3 Figures

We show figures of the correspondence between ROAST's directional and mixed tests and calculation of them using formulae from the Methods section of the manuscript. The alternative hypothesis assumed two

causal transcripts each with log odds ratios of 0.3 and a correlated block of size 20 with correlation 0.7.

S3.1 Under the null hypothesis

Figure S2: Directional test p-values under the null on the log scale

Figure S3: Mixed test p-values under the null on the log scale

S3.2 Under the alternative hypothesis

Figure S4: Mixed test p-values under the null on a linear scale

Figure S5: Directional test p-values under the alternative on the log scale

Figure S6: Mixed test p-values under the alternative on the log scale

Figure S7: Mixed test p-values under the alternative on a linear scale