



Variance component score test for time-course gene set analysis of longitudinal RNA-seq data

Denis Agniel, Boris P. Hejblum

► To cite this version:

Denis Agniel, Boris P. Hejblum. Variance component score test for time-course gene set analysis of longitudinal RNA-seq data. *Biostatistics*, 2017, *Biostatistics*, 18 (4), pp.589-604. 10.1093/biostatistics/kxx005 . hal-01579077

HAL Id: hal-01579077

<https://inria.hal.science/hal-01579077>

Submitted on 6 Oct 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Variance component score test for time-course gene set analysis of longitudinal RNA-seq data

DENIS AGNIEL

Department of Biomedical Informatics, Harvard Medical School, Boston, MA, USA

BORIS P. HEJBLUM *

Department of Biostatistics, Harvard T.H. Chan School of Public Health, Boston, MA, USA

University of Bordeaux, ISPED, INSERM U1219, INRIA SISTM, Bordeaux, FRANCE

Vaccine Research Institute, Créteil, FRANCE

October 3, 2018

Abstract

As gene expression measurement technology is shifting from microarrays to sequencing, the statistical tools available for their analysis must be adapted since RNA-seq data are measured as counts. It has been proposed to model RNA-seq counts as continuous variables using nonparametric regression to account for their inherent heteroscedasticity. In this vein, we propose *tcgsaseq*, a principled, model-free, and efficient method for detecting longitudinal changes in RNA-seq gene sets defined *a priori*. The method identifies those gene sets whose expression varies over time, based on an original variance component score test accounting for both covariates and heteroscedasticity without assuming any specific parametric distribution for the (transformed) counts. We demonstrate that despite the presence of a nonparametric component, our test statistic has a simple form and limiting distribution, and both may be computed quickly. A permutation version of the test is additionally proposed for very small sample sizes. Applied to both simulated data and two real datasets, *tcgsaseq* is shown to exhibit very good statistical properties, with an increase in stability and power when compared to state-of-the-art methods ROAST, edgeR, and DESeq2, which can fail to control the type I error under certain realistic settings. We have made the method available for the community in the R package *tcgsaseq*.

Key words: Gene Set Analysis; Longitudinal data; RNA-seq data; Variance component testing; Heteroscedasticity

*bhejblum@hsph.harvard.edu

1 Introduction

Gene expression is a dynamic biological process of living organisms, whose dysfunction and variation can be related to numerous diseases. During the past two decades, gene expression measurements have developed rapidly, thanks to wide dissemination of microarray technology. In recent years, gene expression measurement technology has been shifting from microarrays to sequencing (RNA-seq) technology. The higher resolution of RNA-seq technology provides a number of advantages over microarrays, among which are the ability to make *de novo* discoveries and an increased sensitivity to low-abundance variants (Marioni and others, 2008; Malone and Oliver, 2011). Indeed, RNA-seq is not restricted to a predefined set of probes like microarrays but can instead measure the genome in its entirety.

A large body of statistical methods have been developed to analyze microarray data. But as technology for measuring gene expression is transitioning to RNA-seq, new methodological challenges arise. RNA-seq produces count data, while microarray analysis techniques generally assume continuity. Due to their underlying count nature, RNA-seq data are intrinsically heteroscedastic. Various approaches have been proposed to deal with these issues, mostly relying on modeling the underlying count nature of the data through the use of Poisson or negative binomial distributions (Marioni and others, 2008; Anders and Huber, 2010; Robinson and others, 2010). Recently, Law and others (2014) have instead proposed to use normal-based methods to analyze RNA-seq data by explicitly modeling the heteroscedasticity and accounting for it by weighting. All of these methods currently make potentially restrictive assumptions about the distribution of the data.

While most methods for gene expression data focus on univariate differential gene expression analysis, it has been shown that gene set analysis (GSA) can be a more powerful and interpretable alternative (Subramanian and others, 2005; Hejblum and others, 2015). GSA uses *a priori* defined gene sets annotated with biological functions and investigates their potential association with biological conditions of interest. There are many different approaches to GSA, and a GSA method is typically defined by the type of hypothesis tested as well as how information across genes is aggregated. There are two primary approaches to specifying tested hypotheses: self-contained hypotheses and competitive hypotheses (Goeman and Bühlmann, 2007). Rahmatallah and others (2016) recently showed that self-contained GSA tests tend to be more powerful and more robust than competitive ones (Goeman and Bühlmann, 2007). Furthermore, some GSA tests rely on univariate gene-level statistics as a first step, aggregating them afterwards in a bottom-up enrichment approach. But when signal strength is weak, single-step, top-down GSA methods relying on direct multivariate modeling are better than those enrichment based ones at leveraging the additional power of GSA over differential gene expression analysis (Hejblum and others, 2015).

As costs keep decreasing for RNA-seq experiments, more complex study designs, such as time-course experiments, have become more common (Dorr

and others, 2015; *Leng and others*, 2015; *Baduel and others*, 2016). However, very few GSA approaches can properly accommodate and test hypotheses more complex than simple differential expression, such as change over time. The ROAST method (*Wu and others*, 2010) is a linear-model-based gene set testing procedure. *Law and others* (2014) have proposed to use it in conjunction with their heteroscedasticity weighting method voom on RNA-seq data, and this combination of voom and ROAST has been identified in the recent review by *Rahmatallah and others* (2016) as one of the top-performing GSA methods for RNA-seq data. Among the eleven methods compared in the review, only three are derived from top-down models incorporating all the genes in the set at once, and ROAST is the only one that is not limited to comparisons of two groups. Additionally, DESeq2 (*Love and others*, 2014) and edgeR (*Robinson and others*, 2010; *McCarthy and others*, 2012), are currently the most prominent approaches used for gene-level differential analysis of RNA-seq data. They both rely on the assumption that gene counts from RNA-seq measurements follow a negative binomial distribution. edgeR can use the ROAST method to propose a built-in framework for self-contained gene set testing, while DESeq2 can only perform gene-wise tests. *Nueda and others* (2014) have also developed a gene-wise method for dealing with time-course RNA-seq data, considering Poisson and negative binomial distributions through generalized linear models.

Another concern, particularly in the longitudinal setting, is adequately accounting for heterogeneity of effects (*Cui and others*, 2016). GSA approaches often overlook the potential longitudinal heterogeneity within a gene set while such heterogeneity is not infrequent (*Ackermann and Strimmer*, 2009) and can be of biological interest (*Prieto and others*, 2006; *Hu and others*, 2013), especially if considered gene sets correspond to biological pathways or networks. *Hejblum and others* (2015) showed the potential statistical power gain when accounting for this heterogeneity in longitudinal microarray studies.

In this article, we propose *tcgsaseq*, a method to analyze RNA-seq data at the gene-set level, with a particular focus on longitudinal studies. We derive a variance component score test, similar to those that have been proposed in other testing situations (*Lin*, 1997; *Wu and others*, 2011; *Huang and Lin*, 2013), that facilitates testing both homogeneous and heterogeneous gene sets simultaneously. Variance component tests offer the speed and simplicity of standard score tests, but potentially gain statistical power by using many fewer degrees of freedom, and have been shown to have locally optimal power in some situations (*Goeman and others*, 2006). Inspired by the voom approach from *Law and others* (2014), we propose to estimate the mean-variance relationship in a more principled way using local linear regression, to account for the inherent heteroscedasticity of the data. Despite this nonparametric step, we demonstrate that the test statistic has a simple limiting distribution that is valid regardless of any model specification. We also propose a permutation version of the test to deal with small sample sizes. Our method is implemented in the R package *tcgsaseq*, available on the Comprehensive R Archive Network at <https://cran.r-project.org/web/packages/tcgsaseq>.

Our general approach to GSA in longitudinal RNA-seq studies has three pri-

primary advantages over existing approaches. First, unlike ROAST, our variance component approach remains valid under model misspecification and does not rely on *ad hoc* aggregation of information across genes. Second, unlike previous approaches to variance component testing in microarray data (Huang and Lin, 2013), our approach can accommodate the intrinsic mean-variance relationship in RNA-seq data, while remaining fast to compute. Third, our test remains powerful even when patients display heterogeneous trajectories over time.

The remainder of the paper is organized as follows. Section 2 describes our variance component score test, while its asymptotic properties are derived in Section 3, and an estimation strategy and practical recommendations are detailed in Section 4. In Section 5, we present numerical studies assessing the potential impact of ignoring the mean-variance relationship in RNA-seq data, and comparing our approach to state-of-the-art methods of gene set analysis. In Section 6, we apply tcgsaseq to the analysis of real data. Final remarks and comments are discussed in Section 7.

2 Variance component score test for longitudinal gene set analysis

2.1 Problem setup

Consider $\mathbf{r}_i = (\mathbf{r}_{i1}^\top, \dots, \mathbf{r}_{iP}^\top)^\top$, $\mathbf{r}_{ij} = (r_{ij1}, \dots, r_{ijn_i})^\top$ a vector of sequence read counts for subject i , that have been mapped to each of P genes measured at times $\mathbf{t}_i = (t_{i1}, \dots, t_{in_i})^\top$. Meanwhile $\mathbf{x}_i = (x_{i1}, \dots, x_{iq})^\top$ is a vector of baseline covariates describing experiment design conditions, all measured on individual i . So the full data considered for analysis are n independent realizations from random vectors $\mathbb{V} = \{\mathbf{V}_i = (\mathbf{r}_i^\top, \mathbf{t}_i^\top, \mathbf{x}_i^\top)^\top\}_{i=1}^n$. Typically, the counts $\mathbf{r}_i$ are normalized in some way in a preprocessing step (Hansen and others, 2012). We take $\mathbf{y}_i$ to be a normalized version of $\mathbf{r}_i$. For example, one standard normalization procedure accounts for the library size $\mathcal{L}_{it} = \sum_{j=1}^P r_{ijt}$ by computing the log-counts per million as

$$y_{ijt} = \log_2 \left(10^6 \times \frac{0.5 + r_{ijt}}{1 + \mathcal{L}_{it}} \right). \quad (1)$$

Based on these data, our interest is in identifying gene sets that have longitudinally changing expression patterns. The rest of this section will develop the proposed test statistic and its limiting distribution.

2.2 The test statistic

We are interested in testing for longitudinal changes in a pre-specified gene set, which for the purposes of illustration we take to be the first p genes, $\mathbf{y}_i = (\mathbf{y}_{i1}, \dots, \mathbf{y}_{ip})$. To develop a variance component score test statistic, we start from the following working model, which is a linear mixed effect model (Laird

and Ware, 1982; Fitzmaurice *and others*, 2012)

$$y_{ijt} = \alpha_{0j} + a_{0ij} + \mathbf{x}_i^\top \boldsymbol{\alpha}_j + \boldsymbol{\phi}_{it}^\top \boldsymbol{\beta}_j + \boldsymbol{\phi}_{it}^\top \boldsymbol{\xi}_{ij} + \epsilon_{ijt} \quad (2)$$

and which can be rewritten as

$$\mathbf{y}_i = \boldsymbol{\alpha}_0 + \mathbf{a}_{0i} + X_i \boldsymbol{\alpha} + \Phi_i \boldsymbol{\beta} + \Phi_i \boldsymbol{\xi}_i + \boldsymbol{\varepsilon}_i \quad (3)$$

where $\boldsymbol{\varepsilon}_i \sim N(0, \Sigma_{\varepsilon i})$, $\mathbf{a}_{0i} \sim N(0, \Sigma_a)$, $\boldsymbol{\alpha}_0$ is a $pn_i \times 1$ vector of gene-specific intercepts α_{0j} , $\mathbf{a}_{0i}$ is a $pn_i \times 1$ vector of random intercepts a_{0ij} , X_i is a $pn_i \times pq$ block diagonal matrix with each block consisting of n_i rows equal to the fixed effect baseline covariates $\mathbf{x}_i$, $\boldsymbol{\alpha}$ is a $pq \times 1$ vector of gene-specific fixed effects $\boldsymbol{\alpha}_j$, $\Phi_i = \Phi(\mathbf{t}_i)$ is a $pn_i \times pK$ block-diagonal matrix encoding the effect of time with each block having t^{th} row $\boldsymbol{\phi}_{it} = \{\phi_1(t_{it}), \dots, \phi_K(t_{it})\}$ for some set of K basis functions $\{\phi_k(\cdot)\}$, $\boldsymbol{\beta}$ is a $pK \times 1$ vector of gene-specific fixed effects of time $\boldsymbol{\beta}_j$, $\boldsymbol{\xi}_i$ is a $pK \times 1$ vector of gene-specific individual-level effects of time, and $\Sigma_{\varepsilon i}$ is a $pn_i \times pn_i$ covariance matrix of measurement errors. Note that $\Sigma_{\varepsilon i}$ is indexed by i and may depend on the mean of $\mathbf{y}_i$. We assume that $\mathbf{a}_{0i} \perp \boldsymbol{\varepsilon}_i$. It is important to note that, in practice, the above model is unlikely to hold. Fortunately the testing procedure we propose is entirely robust to its misspecification.

We are interested in testing the null hypothesis of no longitudinal change in the normalized gene expression for any genes in the gene set

$$H_0 : \text{The mean of } \mathbf{y}_i \text{ does not depend on } \mathbf{t}_i \quad (4)$$

or, that is, $E(\mathbf{y}_i | \mathbf{t}_i) = E(\mathbf{y}_i)$ and $E(\mathbf{y}_i | \mathbf{t}_i, \mathbf{x}_i) = E(\mathbf{y}_i | \mathbf{x}_i)$. Under the model (3), the null hypothesis (4) corresponds to the following null: $H_0 : \boldsymbol{\beta} = \boldsymbol{\xi}_i = \mathbf{0}$, which tests for both homogeneous and heterogeneous dynamics within the gene sets. In section 3, we demonstrate that the corresponding variance component score test can be written as

$$Q = \mathbf{q}^\top \mathbf{q}, \quad \mathbf{q}^\top = n^{-\frac{1}{2}} \sum_{i=1}^n \mathbf{y}_{\mu i}^\top \Sigma_i^{-1} \Phi_i \Sigma_\nu^{1/2} \quad (5)$$

where $\Sigma_i = \Sigma_a + \Sigma_{\varepsilon i}$, $\mathbf{y}_{\mu i} = \mathbf{y}_i - \boldsymbol{\alpha}_0 - X_i \boldsymbol{\alpha}$, Σ_ν is a working covariance matrix defined in (10), and $\Sigma_\nu^{1/2}$ is the symmetric half matrix such that $\Sigma_\nu^{1/2} \Sigma_\nu^{1/2} = \Sigma_\nu$. We also show in the Supplementary Material that, given the dimension of $\mathbf{q}$ (pK) is small relative to the number of individuals (n), the asymptotic distribution of the test statistic is a mixture of χ_1^2 random variables, $Q \rightarrow \sum_{l=1}^{pK} a_l \chi_1^2$, where the mixing coefficients $\{a_l\}$ depend on the covariance of $\mathbf{q}$. This asymptotic distribution holds regardless of the distribution of $\mathbf{y}$.

In the end, p-values may be computed by comparing the observed test statistic Q to the distribution of $\sum_{l=1}^{pK} \hat{a}_l \chi_1^2$ where $\hat{a}_l$ is an estimate of a_l . Details are developed further in Section 3 and the Supplementary Material. When the entries of $\mathbf{q}$ are correlated – i.e. there is correlation between genes in the gene set – then the degrees of freedom for the test based on Q ($\sum_{l=1}^{pK} a_l$) may be much lower than the degrees of freedom for a similar Wald or score test (pK), yielding more power to detect departures from the null hypothesis.

In practice, the parameters α_0 , α , Σ_i , and Σ_ν are unknown and must be estimated. In particular, accounting for the mean-variance relationship encoded in Σ_i is vital in RNA-seq data. Hence we estimate the test statistic as

$$\hat{Q} = \hat{\mathbf{q}}^\top \hat{\mathbf{q}}, \quad \hat{\mathbf{q}}^\top = n^{-\frac{1}{2}} \sum_{i=1}^n (\mathbf{y}_i - \hat{\alpha}_0 - X_i \hat{\alpha})^\top \hat{\Sigma}_i^{-1} \Phi_i \hat{\Sigma}_\nu^{1/2} \quad (6)$$

and we further demonstrate in section 3 that plugging in standard estimates of Σ_ν , α_0 , and α and a nonparametric estimator of Σ_i for the estimated test statistic still yields a similar asymptotic distribution:

$$\hat{Q} \rightarrow \sum_{l=1}^{pK} \tilde{a}_l \chi_1^2 \quad (7)$$

for mixing coefficients $\{\tilde{a}_l\}$ given in the Supplementary Material.

The strength of our approach is that this simple limiting distribution holds even when the model (3) may be misspecified and despite the presence of the nonparametric estimator of Σ_i in (6). In this way, we may account very flexibly for the mean-variance relationship in $\mathbf{y}$ while maintaining a simple, powerful test statistic that does not require any particular model to hold.

The convergence in (7) relies on the central limit theorem. There may of course be situations where the central limit theorem fails to kick in, either when the number of genes in the gene set p is quite large, or else when the number of individuals n is small. We now discuss two modifications for these scenarios.

2.2.1 Testing for homogeneous gene sets

When the number of genes p is large relative to n , the convergence of $\mathbf{q}$ to a limiting normal distribution may be in doubt. In these cases, it might be better to begin from a working model that assumes that all genes share a common trajectory over time as a useful approximation. To wit, we may reduce the number of parameters in the model by taking X_i in model (3) to be a $pn_i \times q$ counterpart of X_i , and similarly α to be $q \times 1$, Φ_i to be $pn_i \times K$, and ξ_i to be $K \times 1$ counterparts, respectively, of the original quantities.

The derivation of the test statistic follows precisely the same lines with updated dimensions. This test will be powerful to detect alternatives where all genes in the set develop in a similar longitudinally. However, this test should be expected to have low power to detect situations where some genes, for example, decline over time while other genes increase over time, as in heterogeneous gene sets.

2.2.2 Using permutation

When n is very small (or alternatively when p is relatively large but the homogeneous strategy described above seems unwise), relying on the limiting distribution (7) may not be accurate. In these cases, permutation may be used

to estimate the empirical distribution of $\widehat{Q}$ under the null (4). To perform the permutations, we simply shuffle the time labels within each individual to get permuted observations for the i th individual $y_{ijt}^* = y_{ij\mathcal{I}_t}$ where $\{\mathcal{I}_1, \dots, \mathcal{I}_{n_i}\}$ is a permutation of $\{1, \dots, n_i\}$. Indeed, under the null, observations of a given gene j for a given individual i are exchangeable, regardless of sampling time. A large number B of permutation-based test statistics $\mathfrak{Q} = \{Q_1^*, \dots, Q_B^*\}$ can thus be generated where each Q_b^* is computed using the b th set of permuted data. P-values may then be computed as:

$$\widehat{p} = \frac{1}{B} \sum_{b=1}^B I\{\widehat{Q} \leq Q_b^*\} \quad (8)$$

3 Properties of the test statistic

In this section, we derive the test statistic and demonstrate its asymptotic distribution assuming all parameters known. We then take up the distribution of the test statistic when estimating all relevant parameters.

3.1 Test statistic derivation

Under the working model (3), $\mathbf{y}_i | \mathbf{a}_{0i}, \boldsymbol{\xi}_i, \mathbf{x}_i, \mathbf{t}_i \sim N(\boldsymbol{\alpha}_0 + \mathbf{a}_{0i} + X_i \boldsymbol{\alpha} + \Phi_i \boldsymbol{\beta} + \Phi_i \boldsymbol{\xi}_i, \Sigma_{\varepsilon_i})$. Integrating over the random intercepts $\mathbf{a}_{0i}$, we can rewrite the model as

$$\mathbf{y}_i | \boldsymbol{\xi}_i, \mathbf{x}_i, \mathbf{t}_i \sim N(\boldsymbol{\mu}_i + \Phi_i \boldsymbol{\theta}_i, \Sigma_i) \quad (9)$$

where $\boldsymbol{\mu}_i = \boldsymbol{\alpha}_0 + X_i \boldsymbol{\alpha}$ denotes time-independent fixed effects, $\boldsymbol{\theta}_i = (\boldsymbol{\beta}_1^\top + \boldsymbol{\xi}_{i1}^\top, \dots, \boldsymbol{\beta}_p^\top + \boldsymbol{\xi}_{ip}^\top)^\top$ denotes combined effects of time.

The test for no longitudinal changes in expression corresponds to the model-based hypothesis $H_0 : \boldsymbol{\theta}_i = \mathbf{0}$. We write $\boldsymbol{\theta}_i$ as $\boldsymbol{\theta}_i = \eta \boldsymbol{\nu}_i$ and we consider the working assumption that the $\{\boldsymbol{\nu}_i\}$ are independently distributed such that $E(\boldsymbol{\nu}_i) = \mathbf{0}$ (under the null) and

$$\text{Var}(\boldsymbol{\nu}_i) = \Sigma_\nu. \quad (10)$$

Under this working assumption, $H_0 : \boldsymbol{\theta}_i = \mathbf{0}$ is equivalent to $H_0 : \eta = 0$. To obtain the variance component test statistic, rewrite the model as: $\mathbf{y}_{\mu i} = \Phi_i \boldsymbol{\theta}_i + \boldsymbol{\varepsilon}_i$ for centered outcome $\mathbf{y}_{\mu i} = \mathbf{y}_i - \boldsymbol{\mu}_i$. Then $\mathbf{y}_{\mu i} | \boldsymbol{\nu}_i \sim N(\Phi_i \eta \boldsymbol{\nu}_i, \Sigma_i)$ and the log-likelihood for $\mathbf{y}_{\mu i}$ can be written as

$$\log \mathcal{L}(\eta) = -\frac{1}{2} \sum_{i=1}^n \left\{ \log |\Sigma_i| + (\mathbf{y}_{\mu i} - \eta \Phi_i \boldsymbol{\nu}_i)^\top \Sigma_i^{-1} (\mathbf{y}_{\mu i} - \eta \Phi_i \boldsymbol{\nu}_i) \right\}. \quad (11)$$

Because the target of inference is η , we marginalize over the nuisance parameter $\boldsymbol{\nu}$ conditional on the observed data to obtain $\mathcal{L}^*(\eta) = E\{\mathcal{L}(\eta) | \mathbb{V}\}$ where the expectation is taken over the distribution of $\boldsymbol{\nu}$. We follow the argument in

Commenges and Andersen (1995) and note that the score at the null value is 0: $\lim_{\eta \rightarrow 0} \partial \log \mathcal{L}^*(\eta) / \partial \eta = E \left(\sum_{i=1}^n \mathbf{y}_{\mu i}^\top \Sigma_i^{-1} \eta \boldsymbol{\nu}_i^\top \boldsymbol{\phi}_i \mid \mathbb{V} \right) = 0$. So we instead consider the score with respect to η^2 , $\lim_{\eta \rightarrow 0} \partial \log \mathcal{L}^*(\eta) / \partial (\eta^2)$, which is equal to

$$E \left\{ \left. \frac{\partial \log \mathcal{L}(\eta)}{\partial \eta} \right|_{\eta=0} \mid \mathbb{V} \right\}^2 + \text{constant} + o_p(1) \quad (12)$$

$$= \left(\sum_{i=1}^n \mathbf{y}_{\mu i}^\top \Sigma_i^{-1} \boldsymbol{\Phi}_i \right) \Sigma_\nu \left(\sum_{i=1}^n \mathbf{y}_{\mu i}^\top \Sigma_i^{-1} \boldsymbol{\Phi}_i \right)^\top + \text{constant} + o_p(1) \quad (13)$$

Thus, after normalizing by n^{-1} , the variance component score test statistic can be written $Q = \mathbf{q}^\top \mathbf{q}$, as described in (5). If the dimension of $\mathbf{q}$ is small relative to the sample size, $\mathbf{q}$ is asymptotically normal by the central limit theorem, which means that the limiting distribution of Q is a mixture of χ_1^2 s (see Supplementary Material for more details).

3.2 Estimated test statistic limiting distribution

In practice, estimates need to be supplied for many of the quantities in the test statistic to be estimated in (6). Following the argument for the limiting distribution of Q , $\hat{Q}$ has a limiting distribution of a mixture of χ_1^2 s so long as $\hat{\mathbf{q}}$ has a limiting normal distribution. To establish this, we define $\tilde{\mathbf{q}} = n^{-\frac{1}{2}} \sum_{i=1}^n (\mathbf{y}_i - \hat{\boldsymbol{\mu}}_i)^\top \mathcal{T}_i$ to be a version of $\hat{\mathbf{q}}$ where $\hat{\Sigma}_i^{-1} \boldsymbol{\Phi}_i \hat{\Sigma}_\nu^{1/2}$ is replaced with its limit, $\mathcal{T}_i = \lim_{n \rightarrow \infty} \hat{\Sigma}_i^{-1} \boldsymbol{\Phi}_i \hat{\Sigma}_\nu^{1/2}$. The quantities $\hat{\mathbf{q}}$ and $\tilde{\mathbf{q}}$ are asymptotically equivalent, and the central limit theorem ensures that $\tilde{\mathbf{q}}$ has a limiting normal distribution. Therefore the asymptotic distribution of $\hat{Q}$ is a mixture of χ_1^2 s (see Supplementary Material for details), and the corresponding testing procedure is valid under any data-generating function and robust to misspecification of the working model 3.

4 Estimation

The test statistic (6) contains the estimated quantities $\hat{\boldsymbol{\alpha}}_0$, $\hat{\boldsymbol{\alpha}}$, $\hat{\Sigma}_\nu$, and $\hat{\Sigma}_i$. In this section, we take up practical issues involved in the estimation of these quantities.

4.1 Estimating model parameters

In this section, we will discuss estimation of $(\hat{\boldsymbol{\alpha}}_0, \hat{\boldsymbol{\alpha}}, \hat{\Sigma}_\nu)$, assuming that we have a consistent estimator of $\hat{\Sigma}_i$ in hand (we leave discussion of estimating $\hat{\Sigma}_i$ to the following section). A natural way to estimate $(\hat{\boldsymbol{\alpha}}_0, \hat{\boldsymbol{\alpha}}, \hat{\Sigma}_\nu)$, taking into account the heteroscedasticity in $\mathbf{y}$, is to fit a weighted mixed effects model corresponding to

$$y_{ijt} = \alpha_{0j} + \mathbf{x}_i^\top \boldsymbol{\alpha}_j + \boldsymbol{\phi}_{it}^\top \boldsymbol{\theta}_{ij} + \tilde{\epsilon}_{ijt}, \quad (14)$$

where the random intercepts a_{0ij} have been integrated out and θ_{ij} is taken to be $N(0, \Sigma_\nu)$. The weights are taken to be $\mathbf{w}_i = \text{diag}(\widehat{\Sigma}_i)^{-1}$. This model corresponds to the model (9).

In practice, fitting the full mixed effects model for each of many gene sets may be computationally demanding. In these cases, the simpler fixed effects model

$$y_{ijt} = \alpha_{0j} + \mathbf{x}_i^\top \boldsymbol{\alpha}_j + \phi_{it}^\top \boldsymbol{\beta}_j + \tilde{\epsilon}_{ijt} \quad (15)$$

may be fit to estimate $\boldsymbol{\alpha}_0$ and $\boldsymbol{\alpha}$. Σ_ν may not be estimated directly from (15). In this case, it may be specified using a working estimate or taken to be the identity matrix. In general, correct specification of Σ_ν may be difficult and may not have much impact on the test statistic. Furthermore, its misspecification does not affect the validity of the test.

4.2 Estimating the mean-variance relationship

One key feature of our approach is to estimate $\text{Cov}(\mathbf{y}_i | \mathbf{x}_i, \mathbf{t}_i, \boldsymbol{\xi}_i) = \Sigma_i$ in such a way as to account for the heteroscedasticity in $\mathbf{y}$. It is important to note that accurate estimation of Σ_i will likely increase power but will not affect the validity or type I error of the testing procedure.

To approximate the mean-variance relationship in $\mathbf{y}$, we use information from all P genes. We assume that the diagonal elements of Σ_i , which we will denote $\mathbf{v} = \{v_{ijt}\} = \{\text{Var}(y_{ijt} | \mathbf{x}_i, \mathbf{t}_i, \boldsymbol{\xi}_i)\}$, may be modeled as a function of their means $\mathbf{m} = \{m_{ijt}\} = \{E(y_{ijt} | \mathbf{x}_i, \mathbf{t}_i, \boldsymbol{\xi}_i)\}$. Specifically, $v_{ijt} = \omega\{m_{ijt}\} + e_{ijt}$ for some unknown function $\omega(\cdot)$ and errors which follow the moment conditions $E(e_{ijt}) = 0, \text{Var}(e_{ijt}) = \tau^2, \tau > 0$.

We follow [Law and others \(2014\)](#) in using local linear regression in estimating $\omega(\cdot)$. We will first write out the form of the estimator if all parameters were known: $\tilde{\omega}_n(x) = \sum_{ijt} \tilde{\ell}_{ijt}(x) v_{ijt}$, $\tilde{\ell}_{ijt}(x) = \frac{\tilde{b}_{ijt}(x)}{\sum_{ijt} \tilde{b}_{ijt}(x)}$, $\tilde{b}_{ijt}(x) = K\{(m_{ijt} - x)/h\} \left\{ \tilde{S}_{n2}(x) - (m_{ijt} - x) \tilde{S}_{n1}(x) \right\}$ where $\tilde{S}_{nd}(x) = \sum_{ijt} K\{(m_{ijt} - x)/h\} (m_{ijt} - x)^d$ for some kernel function $K(\cdot)$ and bandwidth $h > 0$.

The means $\mathbf{m}$ and variances $\mathbf{v}$ could in principle be estimated using some parametric model like (14) for all P genes. Let $(\hat{m}_{ijt}, \hat{v}_{ijt})$ denote the estimated mean and variance for y_{ijt} . Then the mean-variance relationship may be estimated as

$$\hat{\omega}_n(x) = \tilde{\omega}_n(x) |_{m_{ijt}=\hat{m}_{ijt}, v_{ijt}=\hat{v}_{ijt}} \quad (16)$$

As with any smoothing procedure, choice of the bandwidth h is paramount in producing a reliable estimator. Standard cross-validation techniques may be used to select h in practice. Then the diagonal entries of Σ_i can be estimated as $\hat{\omega}_n(\hat{\mathbf{m}}_i)$.

Practical considerations Under the model (14), we could take $\hat{m}_{ijt} = \hat{\alpha}_{0j} + \mathbf{x}_{it}^\top \hat{\alpha}_j + \phi_{it}^\top \hat{\beta}_j + \phi_{it}^\top \hat{\xi}_{ij}$ and $\hat{v}_{ijt} = (y_{ijt} - \hat{\alpha}_{0j} - \mathbf{x}_{it}^\top \hat{\alpha}_j - \phi_{it}^\top \hat{\beta}_j - \phi_{it}^\top \hat{\xi}_{ij})^2$. However, because $P \gg n$, fitting (14) for such a large number of genes is not practical. Instead, the simpler model (15) could be used. Then $\hat{m}_{ijt}$ could be taken to be $\hat{\alpha}_{0j} + \mathbf{x}_{it}^\top \hat{\alpha}_j + \phi_{it}^\top \hat{\beta}_j$ and $\hat{v}_{ijt}$ to be $(y_{ijt} - \hat{\alpha}_{0j} - \mathbf{x}_{it}^\top \hat{\alpha}_j - \phi_{it}^\top \hat{\beta}_j)^2$. A similar approach was used to model heteroscedasticity in the context of linear regression in Carroll (1982).

On the other hand, fitting a local linear regression on all of $\mathbf{m}$ and $\mathbf{v}$ – a total of $P \sum_{i=1}^n n_i$ observations – may be computationally difficult. In order to reduce the number of points used in the nonparametric fit (16), one could follow Law and others (2014) and model the mean-variance relationship at the gene level

$$v_j = \omega(m_j) + e_j. \quad (17)$$

The gene-level mean may be estimated as $\hat{m}_j = n^{-1} \sum_{i=1}^n n_i^{-1} \sum_{t=1}^{n_i} \hat{\alpha}_{0j} + \mathbf{x}_{it}^\top \hat{\alpha}_j + \phi_{it}^\top \hat{\beta}_j$ and the gene-level variance as $\hat{v}_j = n^{-1} \sum_{i=1}^n n_i^{-1} \sum_{t=1}^{n_i} (y_{ijt} - \hat{\alpha}_{0j} - \mathbf{x}_{it}^\top \hat{\alpha}_j - \phi_{it}^\top \hat{\beta}_j)^2$.

5 Numerical study

In this numerical study section, we illustrate both the importance of accounting for RNA-seq data’s heteroscedasticity and the superiority of tcgsaseq in terms of statistical power and robustness. We have performed simulations under two different settings. The first one demonstrates the good behavior of the asymptotic test under various scenarios in synthetic data. The second one focuses on a very realistic situation where sample size is small and gene counts are generated from a distribution of real observed RNA-seq data. Throughout, for the purposes of testing, we take Σ_a and Σ_ν to be the identity for simplicity. In practice, more precise estimation of Σ_a and Σ_ν would serve to increase power. Finally, throughout we use Davies’s approximation method (Davies, 1980) to compute p-values for the mixture of χ^2 s, implemented in the `CompQuadForm` R package (Duchesne and Lafaye De Micheaux, 2010).

5.1 Synthetic data

In this subsection, we use synthetic data to illustrate the behavior of our estimator when the distribution of the data is known. We first look at the importance of accounting for the mean-variance relationship in Σ_i in tcgsaseq, and then we compare tcgsaseq to other competing methods, including ROAST, edgeR-ROAST, and DESeq2. In order to adapt DESeq2 to self-contained gene set testing, we use the minimum p-value test to adequately combine univariate p-values for testing a whole gene set while taking into account gene correlation (Moskvina and Schmidt, 2008; Lin and others, 2011) and we refer to it as DESeq2-min test. We demonstrate that the tcgsaseq testing procedure is

robust to even heavy misspecification of the mean model and the mean-variance relationship, while ROAST, edgeR-ROAST and DESeq2-min test may suffer from extreme lack of power compared to tcgsaseq or inflate type I errors under misspecification.

5.1.1 Mean-variance relationship

To illustrate the importance of estimating the mean-variance relationship, we generated data under the model $y_{ijt} = \log \left\{ \frac{(\mu_{ijt} + 0.5)10^6}{\sum_i \mu_{ijt} + 1} \right\}$ where

$$\mu_{ijt} = \eta_{ijt} \sum_{i,t} \eta_{ijt}/P + b_{ij}t_{it} + \beta t_{it}, \quad \eta_{ijt} = a_{ij} + \alpha_j + \sum_{k=1}^3 x_{itk} + \epsilon_{ijt}, \quad (18)$$

with $\epsilon_{ijt} \sim N(0, \alpha_j^2)$, $\alpha_j \sim \exp(100)$, $a_{ij} \sim N(0, \alpha_j^2/100)$, $x_{itk} \sim N(100, 2500)$, $t_{it} \sim U(0, 1)$, $b_{ij} \sim N(0, 1)$. We set $n = 200, p = 10, P = 1000, n_i = 3$, and β was allowed to be set to 6 different values within a range from 0 to 2: (0, 0.5, 0.75, 1, 1.25, 1.5, 2). We considered three methods of accounting for the mean-variance relationship: (i) gene-level estimates using equation (17); (ii) voom-type estimation [Law and others \(2014\)](#); (iii) specifying Σ_i to be the identity. Throughout this section we test at the 0.05 level. The model (15) was used to generate the test statistic and adjust for covariates, and the gene-level equation (17) was used to estimate the mean-variance relationship. Note that the mean model (15) is misspecified.

When applying our variance component test, type I error was at the nominal level using either strategy (i) or (iii), but was deflated to 0.035 when using strategy (ii). On the other hand, ROAST, which similarly relies on estimating the mean-variance relationship, inflated type I error to 0.56 when using (ii), and had adequate type I errors using the other weighting strategies. Such results indicate that ROAST may drastically inflate type I error rates in some cases, and that using (ii), the voom-type estimator, to account for heteroscedasticity led to worse performance in general.

Figure 1 demonstrates the effect of accounting for heteroscedasticity under the alternative hypothesis. The more accurate modeling of the mean-variance relationship with method (i) yields large power gains over method (ii). Conversely, by drastically mis-modeling the mean-variance relationship, the naive method (iii) yields almost no power for detecting longitudinal changes.

5.2 Comparison to competing methods

In this subsection, we compare tcgsaseq to ROAST, edgeR-ROAST, and DESeq2-min test using a highly misspecified model and using a negative binomial model. For the misspecified model, we again generate data under model (18) with one modification to ensure that the data are positive integers:

$$y_{ijt} = \max \left(\left\lceil \log \left\{ \frac{(\mu_{ijt} + 0.5)10^6}{\sum_i \mu_{ijt} + 1} \right\} \right\rceil, 1 \right). \text{ We set } n = 50, 100, 150 \text{ and } n_i = 5, \text{ and}$$




Figure 1: Power evaluation in synthetic data according to how heteroscedasticity is accounted for, based on 1,000 simulations.

β took values between -2 and 2. Because DESeq2 is computationally intractable at higher sample sizes with large P , we set the total number of genes P to be 100. We used model (17) to compute the mean-variance relationship for both tcgsaseq and ROAST. Gene-based estimates were used for dispersions in DESeq2, and likelihood ratio tests were used to produce test statistics for both edgeR and DESeq2.

The results are depicted in figure 2. We see that tcgsaseq has the highest power at all sample sizes and at all values of β . Though the model is highly misspecified for all methods, the negative-binomial-based methods edgeR-ROAST and DESeq2min test suffer greatly in terms of power, in particular DESeq2min test. In Figure 2, we show ROAST using the mean-variance relationship estimated from model (17), but it's important to note that if the voom-type strategy were used, the type I error for ROAST rises to more than 0.1 at all sample sizes. Asymmetrical power curves around the null are expected for both edgeR and DESeq2 as they both assume the data have a negative binomial distribution (which is asymmetric with a heavier right tail).




Figure 2: Power evaluation in synthetic data comparing tcgsaseq, ROAST, edgeR-ROAST, and DESeq2-min test, based on 1,000 simulations.

Secondly, we generated data under the negative binomial model, a distributional assumption that edgeR and DESeq2 make. The mean of the negative binomial was specified as

$$\mu_{ijt} = \max\{1001 + a_{0i} + x_i + (b_i + \beta + \beta_j)t_{it}x_i, 0\} \quad (19)$$

where $a_{0i} \sim N(0, 1)$, $x_i \sim N(\mu_{xi}, 1)$, $\mu_{xi} \sim \text{exponential}(1/10)$, $b_i \sim N(0, 1)$, $\beta_j \sim N(0, 1)$, and the negative binomial dispersion parameter was set to 1. We again let β vary between -2 and 2 and considered sample sizes of $n = 50, 100, 150$. In this setting, the default procedure to estimate dispersions was used for DESeq2. The results depicted in Figure 3 show that tcgsaseq and ROAST outperform edgeR-ROAST and DESeq2-min test despite the fact that the distributional assumptions of edgeR and DESeq2 are true. The performance of tcgsaseq and ROAST is in general comparable. Furthermore, it is important to highlight that DESeq2-min test does not control the type I error, and the type I error gets worse as sample size increases.




Figure 3: Power evaluation in negative binomial data comparing tcgsaseq, ROAST, edgeR-ROAST, and DESeq2-min test, based on 1,000 simulations.

To reinforce this point, we generated another dataset in which the mean model is correctly specified (which might be more amenable to ROAST), but with an error variance that depends on time. Specifically,

$$y_{ijt} = a_{ij} + 5x_i + \beta_j t_{it} + \epsilon_{ijt} \tau_j t_{it}, \quad (20)$$

where $a_{ij} \sim N(0, 1)$, $x_i \sim N(\mu_i, 1)$, $\mu_i \sim \exp(10)$, $\beta_j = I\{j = 1\}$, $t_{ij} \sim U(0, 1)$, $\epsilon_{ijt} \sim \text{logistic}$, and $\tau_j \sim \exp(10)$. In this case, applying the ROAST method using strategy (ii) again produced a wildly inflated type I error of 0.95, while using method (iii) also produced inflated type I errors of 0.10. On the other hand, the type I error rate was maintained at the nominal level using method (iii) for tcgsaseq but was highly deflated (0.003) using method (i).

5.3 Realistic small samples simulations

In this subsection, we simulated another dataset to illustrate the good behavior of the tcgsaseq permutation test in the realistic setting of a small sample size. We generated data for 6 individuals each measured at 3 time points according to the scenario described in [Law and others \(2014\)](#), using the script provided in their Supplementary Material:

$$m_{ij} = 0.2 + 1/\sqrt{\mu_{ij} + \beta t_{it}} \text{ and } y_{ijt} = \log \left(\frac{(a_{ij} + 0.5)10^6}{\sum_j a_{ij} + 1} \right) \quad (21)$$

where $t_{it} \in \{1, 2, 3\}$, $a_{ij} \sim \text{Pois}(\lambda_{ij})$, $\lambda_{ij} \sim \text{Gamma}(\frac{\kappa_{ij}}{40m_{ij}}, \frac{40m_{ij}}{\kappa_{ij}\mu_{ij}})$, $\kappa_{ij} \sim \chi^2(40)$, and μ_{ij} follows an empirical baseline distribution derived from real RNA-seq counts data, provided in supplementary material of [Law and others \(2014\)](#). This simulation scheme allows us to generate data that realistically resemble real RNA-seq count data. We set $n = 6$ with $n_i = 3$, and gene sets were constructed such that $p \in (30, \dots, 400)$ and for every gene pair (j, j') in the set $\text{cor}(\mu_j, \mu_{j'}) > 0.8$. β was allowed to be set to 8 different values within a range from 0 to 0.5: (0, 0.01, 0.05, 0.1, 0.2, 0.3, 0.4, 0.5).

As shown in Figure 4, under this setting, edgeR-ROAST, ROAST and the proposed permutation test all control the type I error at nominal rate. However DESeq2-min test fails to control the type I error, while it is highly deflated for the asymptotic version of tcgsaseq (0.001) which is not surprising given that with such a small sample size it is unlikely the central limit theorem would have come into effect yet. Nonetheless, we observe a steady and consistent increase of power for our method using the permutation test over the compared state-of-the-art approaches. The deviation from the negative binomial distribution hypothesis (which both edgeR and DESeq2 rely upon), as well as the univariate gene-wise step first needed, explain those poor performances compared to tcgsaseq. Similar results are obtained when adding a random effect of time simulating heterogeneous gene sets (see Supplementary Material). In addition, as gene sets were constructed with correlated genes, these results also indicates that tcgsaseq is robust to and can take advantage of inter-gene correlation.




Figure 4: Power evaluation in realistically simulated data with a small sample size, based on 500 simulations.

6 Application to two real datasets

In this section we present analyses of two real datasets. In both analyses, Σ_a and Σ_v were again taken as the identity while Σ_i was estimated at the gene level using equation (17) for tcgsaseq. The ROAST method was applied in combination with precision weights estimated through the voom approach. Given the sample sizes available in both studies, we used the tcgsaseq permutation test. In both cases, only transformed data were available and not original counts data, preventing us from applying either edgeR-ROAST or DESeq2-min test.

6.1 Longitudinal RNA-seq measurements in successful kidney transplant patients

We analyzed a RNA-seq dataset from [Dorr and others \(2015\)](#), in which gene expression was measured in the peripheral blood mononuclear cells of 32 kidney transplant patients. Gene expression was measured at 4 time points: before transplantation, 1 week after transplantation, 3 months after transplantation and 6 months after transplantation. The patients had no graft rejection at the time of each sample.

We investigated custom gene sets targeted specifically towards kidney transplant. Those kidney-oriented gene sets were derived by the Alberta Transplant Applied Genomics Center from specific pathogenesis-based transcripts ([Halloran and others, 2010](#); [Sellarés and others, 2013](#); [Broin and others, 2014](#)), and their definition is available at <http://atagc.med.ualberta.ca/Research/GeneLists>. We tested for a linear change in gene expression over time while adjusting for patient’s age and gender.

Figure 5 shows the p-values for those 9 gene sets for both tcgsaseq and

ROAST. At a 5% threshold, our approach tcgsaseq identifies three significant gene sets while the combination of voom and ROAST identifies none. Among those significant gene sets, two relate to T-cell gene expression, corroborating the original results from [Dorr and others \(2015\)](#). The gene set annotated as “T-cells” notably includes the gene CD3D, previously highlighted in [Dorr and others \(2015\)](#). On the other hand, gene sets related to transplant rejection such as “Donor-specific antibody” or “Gamma-IFN and Rejection” have much higher p-values here, which is what one would expect as these data include only successful transplant patients. Besides, tcgsaseq also detects a significant change in expression for the gene set related to Mast cells, which have recently been highlighted as playing an ambiguous role in kidney transplant ([Papadimitriou and others, 2013](#)). Mast cells have been linked both to peripheral tolerance ([de Vries and others, 2009](#)), as well as to late graft loss [Jevnikar and Mannon \(2008\)](#). These results both reinforce and broaden the original findings from [Dorr and others \(2015\)](#).




Figure 5: P-values from testing the 9 kidney oriented gene sets investigated

6.2 Time-course RNA-seq comparative study of *Arabidopsis arenosa* physiology

In a recent experiment, [Baduel and others \(2016\)](#) measured time-course gene expression of the plant *Arabidopsis arenosa* through RNA-seq. They sampled 48 plants across 13 weeks at four different time points. [Baduel and others \(2016\)](#) were especially interested in the difference between two populations of *Arabidopsis arenosa*, respectively denoted KA and TBG, that have adopted different flowering strategies. In addition, half of the plants were exposed to cold and short day photoperiods (vernalization) between week 4 and week 10 in order to study the corresponding effect on flowering in both populations. Different plant siblings were sampled at each time point in order to avoid the important stress effect of leaf removal on the plants.

Dealing with this complex experiment design, we used tcgsaseq to address

two separate biological questions: i) which gene sets have a different activation between the two populations, adjusted for the plant age and the cold exposure, ii) which gene sets have a different activation due to cold exposure, adjusted for population differences and plant age. Using two different modeling strategies, we found that one data-driven gene set constructed by [Baduel *and others* \(2016\)](#) from the top 1% differentially expressed genes between the two populations was significant at a 5% threshold for i) and for ii). On the contrary two other data-driven gene sets, again constructed by [Baduel *and others* \(2016\)](#) and identified as the respective population-specific responses to cold exposure, were both significant at this 5% threshold for ii) but not for i). In addition, two gene sets from Gene Ontology associated with salt and cold response pathways, respectively, were also investigated and found significant at the 5% level for ii) but not for i). In comparison, ROAST gave similar results for i), but lacked power for ii) identifying only 1 out of 4 significant gene sets. This analysis corroborates the results obtained by [Baduel *and others* \(2016\)](#), and it further illustrates the good behavior of tcgsaseq in complex time-course RNA-seq studies.

7 Discussion

The proposed method detailed in this article constitutes an innovative and flexible approach for performing gene set analysis of longitudinal RNA-seq gene expression measurements. The approach relies on a principled variance component score test that accounts for the intrinsic heteroscedasticity of RNA-seq data, and for which we derive a simple limiting distribution without requiring any particular model to hold. As illustrated in the previous sections, the good performance of the method when applied to various datasets constitutes a major strength of the method.

Our numerical study shows the importance of accurately accounting for heteroscedasticity when analyzing RNA-seq data. We also demonstrate the robustness of our testing procedure to model misspecification. When comparing our proposed approach to ROAST, edgeR, and DESeq2, state-of-the-art methods, we illustrate the competing methods' sensitivity to model misspecification, failure to control Type I error, and reduced statistical power compared to tcgsaseq. [Goeman *and others* \(2006\)](#) derived a multi-dimensional score-test that has locally optimal power given that i) all of the tested parameters are independent, and ii) there are no nuisance parameters. Their second assumption especially is unrealistic in the context of longitudinal RNA-seq studies targeted by this paper. However, our test being related to theirs, it suggests good performance in statistical power can be expected – just as we demonstrated.

Of particular biological interest in longitudinal studies, the proposed approach can test for both homogeneous and heterogeneous gene sets simultaneously. This is especially relevant for gene sets constructed from biological pathways encompassing both regulators and targets, where different expression dynamics are expected across genes. We also provide a homogeneous version of the testing procedure when heterogeneity of gene dynamics within a gene set

is not of interest. It also results in a power gain compared to GSA methods sensitive only to homogeneous sets. In addition, the proposed solution for accounting for heteroscedasticity can also deal with non-count data and could in principle be applied widely.

As demonstrated by the simulation studies, the proposed approach is also robust to inter-gene correlation within tested gene sets. While in our simulations we do not estimate inter-gene correlations, one could account for them more formally by following [Wang and others \(2009\)](#) and estimating a working correlation matrix with the residuals from an initial gene-wise modeling of $\mathbf{y}$. The resulting estimates could be incorporated into the structure of Σ_i to increase power.

This work features 4 important novelties: i) we propose a statistically sound nonparametric estimation of the inherent mean-variance relationship in RNA-seq data; ii) we present an original variance component score test statistic for which we derive an asymptotic distribution, and we show that it possesses the conventional root- n convergence rate despite the presence of a nonparametric estimator and is robust to model misspecification; iii) we conduct a numerical study comparing the proposed method and illustrating its good statistical performances under various settings; iv) we provide an implementation of the proposed method as an R package available to the community.

These innovations contrast with previous works from either [Law and others \(2014\)](#) or [Hejblum and others \(2015\)](#). While the nonparametric estimation of heteroscedasticity from [Law and others \(2014\)](#) is more of a heuristic procedure, our statistically rigorous approach does not include unmotivated transformations or extrapolation outside the range of the estimating function. Furthermore, we show that our test statistic has the conventional root- n rate despite the presence of the nonparametric estimator. As for [Hejblum and others \(2015\)](#), they focus exclusively on microarray data and therefore completely ignore the heteroscedasticity issue. In addition, while the modeling proposed in this article is similar to theirs, our test statistic is completely different. While they rely on a likelihood ratio test, we propose an original variance component score test that remains valid even if the model was misspecified and is very fast to compute.

Finally, in this paper we focus primarily on longitudinal measurements of RNA-seq. However, our approach directly applies to virtually all RNA-seq study designs, including traditional case-control and more complex studies. Our approach allows researchers to incorporate the natural heteroscedasticity in the data into a powerful test statistic that makes no modeling assumptions. Of note, the inclusion of time-varying covariates would require further assumptions concerning the model to be made to ensure the limiting distribution of the test statistic. Evaluating tcgsaseq’s performance in a broader array of studies is an area for future research.

Software

Software in the form of R code is available on the Comprehensive R Archive Network as an R package `tcgsaseq`.

Acknowledgments

This work was supported by the National Institutes of Health (NIH) [U54 HG007963 to B.P.H.]. The authors express their deepest gratitude to Professor Tianxi Cai, Harvard University, for her help and support in this work. They also thank Pierre Baduel for his help in analyzing the plant physiology data.

Conflict of Interest: None declared.

Supplementary Material

S.1 Properties of the test statistic

In this section, we establish the limiting distribution of the quantity Q and its estimated counterpart $\hat{Q}$.

S.1.1 Limiting distribution of Q

Recall that $Q = \mathbf{q}^\top \mathbf{q}$. Let $\Gamma = \text{Cov}(\mathbf{q})$. Then

$$Q = \mathbf{q}^\top \Gamma^{-1/2} \Gamma \Gamma^{-1/2} \mathbf{q} \quad (22)$$

$$= \mathbf{u}^\top U A U^\top \mathbf{u} + o_p(1) \quad (23)$$

where U is an orthonormal matrix of the eigenvectors of Γ , A is a diagonal matrix of the eigenvalues of Γ , and $\mathbf{u} = \Gamma^{-1/2} \mathbf{q}$ is asymptotically standard multivariate normal by the central limit theorem. Noting that, because U is orthonormal, $U^\top \mathbf{u}$ is also asymptotically standard normal, $\mathbf{u}^\top U A U^\top \mathbf{u} = \sum_{k=1}^{pK} a_k (u_k^*)^2$ where u_k^* is an element of the asymptotically standard normal $U^\top \mathbf{u}$ and a_k an eigenvalue of Γ . It finally follows that $Q \sim \sum_{l=1}^{pK} a_l \chi_1^2$.

S.1.2 Limiting distribution of $\hat{Q}$

Recall that $\hat{Q} = \hat{\mathbf{q}}^\top \hat{\mathbf{q}}$. We first show that $\hat{\mathbf{q}}$ and $\tilde{\mathbf{q}}$ are asymptotically equivalent.

$$\hat{\mathbf{q}}^\top - \tilde{\mathbf{q}}^\top = n^{-\frac{1}{2}} \sum_{i=1}^n (\mathbf{y}_i - \hat{\boldsymbol{\mu}}_i)^\top \hat{\Sigma}_i^{-1} \Phi_i \hat{\Sigma}_i^{1/2} - n^{-\frac{1}{2}} \sum_{i=1}^n (\mathbf{y}_i - \hat{\boldsymbol{\mu}}_i)^\top \mathcal{T}_i \quad (24)$$

$$= n^{-\frac{1}{2}} \sum_{i=1}^n (\mathbf{y}_i - \hat{\boldsymbol{\mu}}_i)^\top (\hat{\mathcal{T}}(\tilde{\mathbf{x}}_i) - \mathcal{T}(\tilde{\mathbf{x}}_i)) \quad (25)$$

for the processes $\mathcal{T}(\tilde{\mathbf{x}}_i) = \Sigma(\tilde{\mathbf{x}}_i)^{-1}\Phi(\tilde{\mathbf{x}}_i)\Sigma_\nu^{1/2} = \Sigma_i^{-1}\Phi_i\Sigma_\nu^{1/2}$ and $\hat{\mathcal{T}}(\tilde{\mathbf{x}}_i) = \hat{\Sigma}(\tilde{\mathbf{x}}_i)^{-1}\Phi(\tilde{\mathbf{x}}_i)\hat{\Sigma}_\nu^{1/2} = \hat{\Sigma}_i^{-1}\Phi_i\hat{\Sigma}_\nu^{1/2}$ as a function of $\tilde{\mathbf{x}}_i = (\mathbf{x}_i, \mathbf{t}_i)$ where $\Sigma(\tilde{\mathbf{x}}_i) = \Sigma_i$, $\Phi(\tilde{\mathbf{x}}_i) = \Phi_i$ and $\hat{\Sigma}(\tilde{\mathbf{x}}_i) = \hat{\Sigma}_i$. Here we use the fact that Σ_i only depends on $\tilde{\mathbf{x}}_i$.

We can thus write

$$\hat{\mathbf{q}} - \tilde{\mathbf{q}} = n^{-\frac{1}{2}} \sum_{i=1}^n \int (\mathbf{y}_i - \hat{\boldsymbol{\mu}}_i)^\top (\hat{\mathcal{T}}(\tilde{\mathbf{x}}) - \mathcal{T}(\tilde{\mathbf{x}})) d\{I\{\tilde{\mathbf{x}}_i \leq \tilde{\mathbf{x}}\}\} \quad (26)$$

$$= \int (\hat{\mathcal{T}}(\tilde{\mathbf{x}}) - \mathcal{T}(\tilde{\mathbf{x}}))^\top d \left\{ n^{-\frac{1}{2}} \sum_{i=1}^n (\mathbf{y}_i - \hat{\boldsymbol{\mu}}_i) I\{\tilde{\mathbf{x}}_i \leq \tilde{\mathbf{x}}\} \right\} \quad (27)$$

Now since $E(\mathbf{y}_i - \hat{\boldsymbol{\mu}}_i | \tilde{\mathbf{x}}_i) = 0$ under (2.4), $n^{-\frac{1}{2}} \sum_{i=1}^n (\mathbf{y}_i - \hat{\boldsymbol{\mu}}_i) I\{\tilde{\mathbf{x}}_i \leq \tilde{\mathbf{x}}\}$ converges to a Gaussian process in $\tilde{\mathbf{x}}$. Using that fact and $\hat{\mathcal{T}}(\tilde{\mathbf{x}}) \rightarrow_p \mathcal{T}(\tilde{\mathbf{x}})$ it can be shown that

$$\int (\hat{\mathcal{T}}(\tilde{\mathbf{x}}) - \mathcal{T}(\tilde{\mathbf{x}}))^\top d \left\{ n^{-\frac{1}{2}} \sum_{i=1}^n (\mathbf{y}_i - \hat{\boldsymbol{\mu}}_i) I\{\tilde{\mathbf{x}}_i \leq \tilde{\mathbf{x}}\} \right\} = o_p(1). \quad (28)$$

This means that $\hat{\mathbf{q}}$ and $\tilde{\mathbf{q}}$ are asymptotically equivalent, and we can examine the limiting behavior of $\tilde{\mathbf{q}}$ to understand the limiting behavior of $\hat{\mathbf{q}}$.

Now, ignoring the intercept $\boldsymbol{\alpha}_0$ for notational simplicity,

$$\tilde{\mathbf{q}} = n^{-\frac{1}{2}} \sum_{i=1}^n (\mathbf{y}_i - \hat{\boldsymbol{\mu}}_i)^\top \mathcal{T}_i \quad (29)$$

$$= n^{-\frac{1}{2}} \sum_{i=1}^n (\mathbf{y}_i - X_i \boldsymbol{\alpha})^\top \mathcal{T}_i - n^{-\frac{1}{2}} \sum_{i=1}^n (\hat{\boldsymbol{\alpha}} - \boldsymbol{\alpha})^\top X_i^\top \mathcal{T}_i \quad (30)$$

$$= n^{-\frac{1}{2}} \sum_{i=1}^n (\mathbf{y}_i - X_i \boldsymbol{\alpha})^\top \mathcal{T}_i - n^{-\frac{1}{2}} \sum_{i=1}^n \mathcal{U}_i(\boldsymbol{\alpha})^\top n^{-1} \sum_{i=1}^n (X_i^\top \mathcal{T}_i) + o_p(1) \quad (31)$$

$$= n^{-\frac{1}{2}} \sum_{i=1}^n \{(\mathbf{y}_i - X_i \boldsymbol{\alpha})^\top \mathcal{T}_i - \mathcal{U}_i(\boldsymbol{\alpha})^\top E(X_i^\top \mathcal{T}_i)\} + o_p(1) \quad (32)$$

$$= n^{-\frac{1}{2}} \sum_{i=1}^n \mathcal{Q}_i + o_p(1) \quad (33)$$

where $\mathcal{U}_i(\boldsymbol{\alpha})$ corresponds to the influence function for $\hat{\boldsymbol{\alpha}}$ from a null parametric model and $\mathcal{Q}_i = (\mathbf{y}_i - X_i \boldsymbol{\alpha})^\top \mathcal{T}_i - \mathcal{U}_i(\boldsymbol{\alpha})^\top E(X_i^\top \mathcal{T}_i)$. Thus, clearly $\tilde{\mathbf{q}}$ has a limiting normal distribution by the central limit theorem, which further suggests that the asymptotic distribution of $\hat{\mathbf{q}}$ is a mixture of χ_1^2 s where the mixing probabilities are given by the eigenvalues of $\text{Cov}(\mathcal{Q}_i)$.

S.2 The voom procedure

Let r_{ijt} be the read count from individual i for gene j at time t . $L_i = \sum_{j=1}^G r_{ijt}$ is then the library size, that can vary from one sample to another. [Law and](#)

[others \(2014\)](#) propose the following transformation to obtain log-counts per million value from sample i for gene g :

$$y_{ijt} = \log_2 \left(10^6 \frac{0.5 + r_{ijt}}{1 + L_i} \right) \quad (34)$$

Counts are offset away from zero by 0.5 to avoid taking the log of zero (and this also reduces the variability of log-cpm for low expression genes) while the library size is offset by 1 to ensure that any log-cpm is below 1 million. Let X be the design matrix of factors that are suspected to influence gene expression.

The voom procedure to estimate mean-variance weights proposed by [Law and others \(2014\)](#) is the following:

1. $\hat{\alpha}$ is the OLS estimate from the linear model: $y_{ijt} = \mathbf{x}_i^T \alpha + \varepsilon_{ijt}$ with $\varepsilon_{ijt} \sim \mathcal{N}(0, \sigma_j)$
2. Let $\hat{s}_j = \sqrt{\sum_{i=1}^n (y_{ijt} - \mathbf{x}_i^T \hat{\alpha})^2}$
3. Let $\tilde{r}_j = \frac{1}{n} \sum_{i=1}^n y_{ijt} + \log_2 \left(\prod_{i=1}^n (1 + \sum_{g=1}^G r_{ijt}) \right)^{1/n} - \log_2(10^6)$ be the average log-count value.
4. $\hat{f}(\cdot)$ is the predictor obtained from the LOWESS regression ([Cleveland, 1979](#)) of $\sqrt{\hat{s}_j}$ over $\tilde{r}_j$. The final precision weights are then defined as follows:

$$w_{ijt} = \left[\hat{f} \left(\mathbf{x}_i^T \hat{\alpha} + \log_2 \left(1 + \sum_{g=1}^G r_{ijt} \right) - \log_2(10^6) \right) \right]^{-4} \quad (35)$$

S.3 Additional simulations

Here we provide additional simulation results in the realistic setting of a small sample size. We generated data for 6 individuals, each measured at 3 time points according to the scenario described in [Law and others \(2014\)](#), using the script provided in their Supplementary Material:

$$y_{ijt} = \log \left(\frac{(a_{ij} + 0.5)10^6}{\sum_j a_{ij} + 1} \right) + \gamma_j t_{it} + \epsilon_{ijt} \quad (36)$$

where $t_{it} \in \{1, 2, 3\}$, $\epsilon_{ijt} \sim N(0, 0.05)$, $a_{ij} \sim \text{Pois}(\lambda_{ij})$, $\lambda_{ij} \sim \text{Gamma}(\frac{\kappa_{ij}}{40m_{ij}}, \frac{40m_{ij}}{\kappa_{ij}\mu_{ij}})$, $\kappa_{ij} \sim \chi^2(40)$, $m_{ij} = 0.2 + 1/\sqrt{\mu_{ij}}$, $\gamma_j \sim N(0, \sigma_\gamma)$, and μ_{ij} follows an empirical baseline distribution derived from real data, provided in supplementary material of [Law and others \(2014\)](#). We set $n = 18$ with $n_i = 3$, and gene sets were constructed such that $p \in (30, \dots, 400)$ and for every gene pair (j, j') in the set $\text{cor}(\mathbf{a}_j, \mathbf{a}_{j'}) > 0.8$. σ_γ varied among $(0, 0.02, 0.05, 0.1, 0.15, 0.2, 0.25, 0.3)$. The results are very similar to those for homogeneous gene sets, see Figure 6.




Figure 6: Power evaluation for heterogeneous gene sets in realistically simulated data with a small sample size, based on 500 simulations.

References

- ACKERMANN, MARIT AND STRIMMER, KORBINIAN. (2009). A general modular framework for gene set enrichment analysis. *BMC Bioinformatics* **10**(1), 47.
- ANDERS, SIMON AND HUBER, WOLFGANG. (2010). Differential expression analysis for sequence count data. *Genome Biology* **11**(10), R106.
- BADUEL, PIERRE, ARNOLD, BRIAN *and others*. (2016). Habitat-Associated Life History and Stress-Tolerance Variation in *Arabidopsis arenosa*. *Plant Physiology* **171**(1), 437–451.
- BROIN, P. O., HAYDE, N. *and others*. (2014). A pathogenesis-based transcript signature in donor-specific antibody-positive kidney transplant patients with normal biopsies. *Genomics Data* **2**, 357–360.
- CARROLL, R.J. (1982). Adapting for heteroscedasticity in linear models. *The Annals of Statistics* **10**(4), 1224–1233.

- CLEVELAND, WILLIAM S. (1979). Robust Locally Weighted Regression and Smoothing Scatterplots. *Journal of the American Statistical Association* **74**(368), 829–836.
- COMMENGES, DANIEL AND ANDERSEN, P K. (1995). Score Test of Homogeneity for Survival Data. *Lifetime data analysis* **1**, 145–156.
- CUI, SHIQI, JI, TIEMING, LI, JILONG, CHENG, JIANLIN AND QIU, JING. (2016). What if we ignore the random effects when analyzing RNA-seq data in a multifactor experiment. *Statistical Applications in Genetics and Molecular Biology* **15**(2), 87–105.
- DAVIES, ROBERT B. (1980). Algorithm AS 155: The Distribution of a Linear Combination of chi-2 Random Variables. *Journal of the Royal Statistical Society. Series C (Applied Statistics)* **29**(3), 323–333.
- DE VRIES, VICTOR C, PINO-LAGOS, KARINA *and others*. (2009). The enigmatic role of mast cells in dominant tolerance. *Current Opinion in Organ Transplantation* **14**(4), 332–337.
- DORR, CASEY, WU, BAOLIN *and others*. (2015). Differentially expressed gene transcripts using RNA sequencing from the blood of immunosuppressed kidney allograft recipients. *PLoS ONE* **10**(5), 1–14.
- DUCHESNE, PIERRE AND LAFAYE DE MICHEAUX, PIERRE. (2010). Computing the distribution of quadratic forms: Further comparisons between the Liu-Tang-Zhang approximation and exact methods. *Computational Statistics and Data Analysis* **54**(4), 858–862.
- FITZMAURICE, GARRETT M, LAIRD, NAN M AND WARE, JAMES H. (2012). *Applied longitudinal analysis*, Volume 998. John Wiley & Sons.
- GOEMAN, JELLE J. AND BÜHLMANN, PETER. (2007). Analyzing gene expression data in terms of gene sets: Methodological issues. *Bioinformatics* **23**(8), 980–987.
- GOEMAN, J J, VAN DE GEER, S A AND VAN HOUWELINGEN, H C. (2006). Testing against a high dimensional alternative. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **68**, 477–493.
- HALLORAN, P. F., DE FREITAS, D. G. *and others*. (2010). The molecular phenotype of kidney transplants: Personal viewpoint. *American Journal of Transplantation* **10**(10), 2215–2222.
- HANSEN, KASPER D, IRIZARRY, RAFAEL A AND WU, ZHIJIN. (2012). Removing technical variability in RNA-seq data using conditional quantile normalization. *Biostatistics* **13**(2), 204–216.
- HEJBLUM, BORIS P., SKINNER, JASON AND THIÉBAUT, RODOLPHE. (2015). Time-Course Gene Set Analysis for Longitudinal Gene Expression Data. *PLOS Computational Biology* **11**(6), e1004310.

- HU, YUXUAN, GAO, LIN *and others*. (2013). Detection of deregulated modules using deregulatory linked path. *PloS one* **8**(7), e70412.
- HUANG, YEN-TSUNG AND LIN, XIHONG. (2013). Gene set analysis using variance component tests. *BMC Bioinformatics* **14**(1), 210.
- JEVNIKAR, ANTHONY M. AND MANNON, ROSLYN B. (2008). Late kidney allograft loss: What we know about it, and what we can do about it. *Clinical Journal of the American Society of Nephrology* **3**(SUPPL. 2), 56–67.
- LAIRD, NAN M AND WARE, JAMES H. (1982). Random-effects models for longitudinal data. *Biometrics* **38**(4), 963–974.
- LAW, CHARITY W, CHEN, YUNSHUN *and others*. (2014). voom: Precision weights unlock linear model analysis tools for RNA-seq read counts. *Genome biology* **15**(2), R29.
- LENG, NING, LI, YUAN *and others*. (2015). EBSeq-HMM: A Bayesian approach for identifying gene-expression changes in ordered RNA-seq experiments. *Bioinformatics* **31**(16), 2614–2622.
- LIN, XIHONG. (1997). Variance component testing in generalised linear models with random effects. *Biometrika* **84**(2), 309–326.
- LIN, XINYI, CAI, TIANXI *and others*. (2011). Kernel machine SNP-set analysis for censored survival outcomes in genome-wide association studies. *Genetic Epidemiology* **35**(7), 620–631.
- LOVE, MICHAEL I, HUBER, WOLFGANG AND ANDERS, SIMON. (2014). Moderated estimation of fold change and dispersion for RNA-seq data with DESeq2. *Genome biology* **15**(12), 550.
- MALONE, JOHN H AND OLIVER, BRIAN. (2011). Microarrays, deep sequencing and the true measure of the transcriptome. *BMC Biology* **9**(1), 34.
- MARIONI, JOHN C, MASON, CHRISTOPHER E *and others*. (2008). RNA-seq: An assessment of technical reproducibility and comparison with gene expression arrays. *Genome Research* **18**(9), 1509–1517.
- MCCARTHY, DAVIS J., CHEN, YUNSHUN AND SMYTH, GORDON K. (2012). Differential expression analysis of multifactor RNA-Seq experiments with respect to biological variation. *Nucleic Acids Research* **40**(10), 4288–4297.
- MOSKVINA, VALENTINA AND SCHMIDT, KARL MICHAEL. (2008). On multiple-testing correction in genome-wide association studies. *Genetic Epidemiology* **32**(6), 567–573.
- NUEDA, MARA JOS, TARAZONA, SONIA AND CONESA, ANA. (2014). Next maSigPro: updating maSigPro bioconductor package for RNA-seq time series. *Bioinformatics* **30**(18), 2598–602.

- PAPADIMITRIOU, J. C., DRACHENBERG, C. B. *and others.* (2013). Mast cell quantitation in renal transplant biopsy specimens as a potential marker for the cumulative burden of tissue injury. *Transplantation Proceedings* **45**(4), 1469–1471.
- PRIETO, C., RIVAS, M. J. *and others.* (2006, 5). Algorithm to find gene expression profiles of deregulation and identify families of disease-altered genes. *Bioinformatics* **22**(9), 1103–10.
- RAHMATALLAH, YASIR, EMMERT-STREIB, FRANK AND GLAZKO, GALINA. (2016). Gene set analysis approaches for RNA-seq data: performance evaluation and application guideline. *Briefings in Bioinformatics* **17**(3), 393–407.
- ROBINSON, MARK D., MCCARTHY, DAVIS J. AND SMYTH, GORDON K. (2010). edgeR: a Bioconductor package for differential expression analysis of digital gene expression data. *Bioinformatics* **26**(1), 139–140.
- SELLARÉS, J., REEVE, J. *and others.* (2013). Molecular diagnosis of antibody-mediated rejection in human kidney transplants. *American Journal of Transplantation* **13**(4), 971–983.
- SUBRAMANIAN, ARAVIND, TAMAYO, PABLO *and others.* (2005). Gene set enrichment analysis: a knowledge-based approach for interpreting genome-wide expression profiles. *Proceedings of the National Academy of Sciences of the United States of America* **102**(43), 15545–15550.
- WANG, LILY, CHEN, XI *and others.* (2009). A unified mixed effects model for gene set analysis of time course microarray experiments. *Statistical Applications in Genetics and Molecular Biology* **8**(1), 1–18.
- WU, DI, LIM, ELGENE *and others.* (2010). ROAST: Rotation gene set tests for complex microarray experiments. *Bioinformatics* **26**(17), 2176–2182.
- WU, MICHAEL C., LEE, SEUNGGEUN *and others.* (2011). Rare-variant association testing for sequencing data with the sequence kernel association test. *American Journal of Human Genetics* **89**(1), 82–93.