

Benchmarking association analyses of continuous exposures with RNA-seq in observational studies

Tamar Sofer, Nuzulul Kurniansyah, François Aguet, Kristin Ardlie, Peter Durda, Deborah A. Nickerson, Joshua D. Smith, Yongmei Liu, Sina A. Gharib, Susan Redline, Stephen S. Rich, Jerome I. Rotter and Kent D. Taylor

Corresponding author: Tamar Sofer, Department of Medicine, Harvard Medical School, Tamar Sofer, 221 Longwood Ave, Boston, MA 02115, USA, Tel.: +1-857-307-0347; Fax: +1-617-278-6946; E-mail: tsofer@bwh.harvard.edu

†These authors contributed equally to the work.

Abstract

Large datasets of hundreds to thousands of individuals measuring RNA-seq in observational studies are becoming available. Many popular software packages for analysis of RNA-seq data were constructed to study differences in expression signatures in an experimental design with well-defined conditions (exposures). In contrast, observational studies may have varying levels of confounding transcript-exposure associations; further, exposure measures may vary from discrete (exposed, yes/no) to continuous (levels of exposure), with non-normal distributions of exposure. We compare popular software for gene expression—DESeq2, edgeR and limma—as well as linear regression-based analyses for studying the association of continuous exposures with RNA-seq. We developed a computation pipeline that includes transformation, filtering and generation of empirical null distribution of association *P*-values, and we apply the pipeline to compute empirical *P*-values with multiple testing correction. We employ a resampling approach that allows for assessment of false positive detection across methods, power comparison and the computation of quantile empirical *P*-values. The results suggest that linear regression methods are substantially faster with better control of false detections than other methods, even with the resampling method to compute empirical *P*-values. We provide the proposed pipeline with fast algorithms in an R package Olivia, and implemented it to study the associations of measures of sleep disordered breathing with RNA-seq in peripheral blood mononuclear cells in participants from the Multi-Ethnic Study of Atherosclerosis.

Key words: observational studies; continuous exposure; non-normality; RNA-seq; empirical *P*-values

Tamar Sofer is an assistant professor at the Harvard University and the director of the Biostatistics and Genomics core at the Program of Sleep Medicine Epidemiology at the Brigham and Women's Hospital.

Nuzulul Kurniansyah is a bioinformatician at the Program of Sleep Medicine Epidemiology at the Brigham and Women's Hospital.

François Aguet is a computational biologist at the Broad Institute of MIT and Harvard.

Kristin Ardlie is an institute scientist at the Broad Institute of MIT and Harvard.

Peter Durda is a faculty scientist at the Department of Pathology & Laboratory Medicine at the University of Vermont.

Deborah A. Nickerson is a professor of Genome Sciences and the Principal Investigator of the Human Genetics and Translational Genomics program at the University of Washington.

Joshua D. Smith is the director of the Bioinformatics of the Human Genetics and Translational Genomics program at the University of Washington.

Yongmei Liu is a professor at the Duke Molecular Physiology Institute.

Sina A. Gharib is a professor and the director of the Computational Medicine Core at the Center of Lung Biology at the University of Washington.

Susan Redline is a professor at the Harvard School of Medicine and the director of the program in Sleep Medicine Epidemiology at Brigham and Women's Hospital.

Stephen S. Rich is a professor and the director of the center for Public Health Genomics at the University of Virginia.

Jerome I. Rotter is a professor and the director of the Institute for Translational Genomics and Population Sciences at the Harbor-UCLA Medical Center at the Lundquist Institute.

Kent D. Taylor is a professor and an investigator at the Harbor-UCLA Medical Center at the Lundquist Institute.

Submitted: 15 February 2021; Received (in revised form): 12 April 2021

Introduction

Many studies of phenotypes associated with gene expression from RNA-seq consist of small sample sizes (tens of subjects) and are focused on comparisons of transcriptional expression patterns between well-delineated states, such as different experimental conditions, tumor versus non-tumor cells [1, 2] and disease versus non-disease groups [3]. Some studies are designed to identify differential expression across hidden, discrete conditions [4]. Epidemiological cohorts have recently utilized stored samples to facilitate the use of RNA-seq data in studies of association with subclinical phenotypes such as blood biomarkers, imaging and other physiological measures, with often continuous measures being used in statistical analyses.

High-throughput RNA sequencing enables broad assaying of a sample's transcriptome [5] and has been in increasing use for over a decade [6]. A large variety of analytic and statistical approaches have been developed to address scientific questions such as alternative splicing, differential expression, and more [4, 7–11], often building on methods developed for analyses of expression microarrays [12–14]; comprehensive reviews are available [15–19]. In addition, a large body of work is being developed for single-cell RNA-seq [20]. In this work, we are specifically interested in differential expression analysis with continuous exposures, and we assume that count data are already prepared and available to the analyst. Popular software packages for differential expression analysis include the DESeq2 R package [9], which models the expression counts as following a negative binomial distribution, with shrinkage imposed on both the mean and the dispersion parameters, based on estimates from the entire transcriptome, or user-supplied values. EdgeR [7] uses a negative binomial model similar to the DESeq2 model for transcript counts, in combination with overdispersion moderation. EdgeR was primarily designed for differential expression analysis between two groups when at least one of the groups has replicated measurements [21]. Limma [22] uses linear models, which are very flexible and can effectively accommodate many study designs and hypotheses. Similar to the DESeq2 and edgeR packages, Limma also uses an empirical Bayes method to borrow information across transcripts to estimate a global variance parameter that is applied for the computation of variance parameters of each single transcript. It uses log transformation and weighting, known as the 'voom' transformation, in the final linear model that is used for differential expression analysis. We refer to it henceforth as the limma-voom. Prior to differential expression analysis, library normalization is performed [23]. Popular approaches are the TMM (trimmed-means of M-values) normalization [24], implemented in edgeR and the size factors normalization [25], implemented in DESeq2.

Sleep disordered breathing (SDB) phenotypes, such as the Apnea-Hypopnea Index (AHI), the number of apnea and hypopnea events per hour of sleep, provides a quantitative assessment of the severity of the disorder, with no clear threshold above which different biological processes occur (although thresholds are used for clinical decision-making and health insurance reimbursement). Association analysis with continuous exposures provides different challenges than those traditionally encountered. The distribution of such exposures may have strong effects on the association analysis results, regardless of the underlying associations, due to the combination of skewed exposure distributions and the distribution of RNA-seq read count data, which are generally over-dispersed with occasional extreme values. As observational study data analyses may include covariates, statistical methods from experimental studies (e.g. exact tests) cannot be applied.

In this article, we compare the DESeq2, edgeR and limma-voom analysis approaches for differential expression analysis, with linear regression-based approaches that do not use the empirical Bayes approach for estimating variance parameters across the transcriptome. We study the computation of *P*-values using resampling of phenotype residuals, while preserving the structure of the data. This addresses the limitation of permutation noted by others in the context of differential expression analysis of RNA-seq [22], where permutation may not be tuned to test a specific null hypothesis because in its standard form it 'breaks' all relationships between the permuted variable and the rest of the dataset. Finally, we study the use of empirical *P*-values that tune the original *P*-values based on the residual resampling scheme. Throughout, we use a dataset with SDB phenotypes and RNA-seq from the Multi-Ethnic Study of Atherosclerosis (MESA) as a case study. We demonstrate the statistical implications of performing association analysis of RNA-seq with continuous, non-normal exposures, compare analysis methods and develop recommendations.

Methods

The Multi-Ethnic Study of Atherosclerosis

MESA is a longitudinal cohort study, established in 2000, that prospectively collected risk factors for development of subclinical and clinical cardiovascular disease among participants in six field centers across the United States (Baltimore City and Baltimore County, MD; Chicago, IL; Forsyth County, NC; Los Angeles County, CA; Northern Manhattan and the Bronx, NY; and St. Paul, MN). RNA was extracted from whole blood drawn in Exam 5 (2010–2012). RNA-seq was measured via the Trans-Omics in Precision Medicine (TOPMed) program. The current analysis considers $N = 462$ individuals who participated in an ancillary sleep study shortly after Exam 5, in 2010–2013 [26, 27]. Sleep data were collected using standardized full in-home level-2 polysomnography (Compumedics Somte Systems, Abbotsville, Australia), as described before [27]. Of the 462 participants in the current analysis, there were 98 African-Americans (AA), 200 European-Americans (EA) and 164 Hispanic-Americans (HA). RNA sequencing in MESA is briefly described in the Supplementary Materials available online at <https://academic.oup.com/bib>. See the Availability section for information about obtaining the MESA dataset.

SDB measures

As examples for continuous exposures from population-based studies, we took three SDB measures: (i) the AHI, defined as the number of apnea (breathing cessation) and hypopnea (at least 30% reduction of breath volume, accompanied by 3% or higher reduction of oxyhemoglobin saturation compared with the baseline saturation) per 1 h of sleep; (ii) minimum oxyhemoglobin saturation during sleep (MinO2) and (iii) average oxyhemoglobin saturation during sleep (AvgO2). We chose these traits because they are clinically relevant, often used in sleep research studies, and represent exposures that may alter gene expression (via hypoxemia and sympathetic activation). The AHI had the least skewed distribution of the considered phenotypes, and AvgO2 had the longest 'tail' of small values in the residual distribution. Residuals were obtained by regression the sleep measures on age, sex, body mass index (BMI), study center and self-reported race/ethnic group.

Compared tests of associations between exposure and transcripts

We compared the standard packages DESeq2, edgeR, limma and linear regression-based approaches, in which we always applied log transformation on the transcript counts, and then applied linear regression. Because some of the observed transcript count values are zero, which cannot be log transformed, we compared a few approaches for replacing zero values. For a given transcript j , denote the minimum observed transcript level that is higher than zero by $m_j = \min\{t_{j1}, \dots, t_{jn} : t_{ji} > 0 \text{ for } i = 1, \dots, n\}$. We compared the following approaches, applied on each transcript $j, j = 1, \dots, k$ separately:

1. SubHalfMin: Replace zero values with $\frac{m_j}{2}$.
2. AddHalfMin: Replace all values t_{ij} by $t_{ij} + \frac{m_j}{2}$.
3. AddHalf: Replace all values t_{ij} by $t_{ij} + \frac{1}{2}$.

Conceptual framework for studying analysis approaches

To study performance of various analysis approaches, we performed simulation studies. Simulation study 1 was used to assess type 1 error across methods when using output P -values, and when using ‘empirical P -values’, which are P -values that account for true distribution of the P -values under the null hypothesis of no association between the exposure and RNA-seq and are described later. Simulation study 2 was used to assess power in transcriptome-wide analysis settings, when using methods that control the type 1 error according to simulation study 1. In addition, we performed a simulation study (see Supplementary Materials available online at <https://academic.oup.com/bib>) to assess power for testing of individual transcript according to various distributional characteristics of transcript counts. The goal was to identify approaches for filtering transcripts for association analysis that will optimize power. All simulations used a ‘residual permutation’ approach (below). The reported criteria for declaring differentially expressed transcripts were false discovery rate (FDR) controlling P -values < 0.05 based on the Benjamini–Hochberg (BH) procedure, and based on the local FDR procedure implemented in the qvalue R package, family-wise error rate (FWER) controlling P -values < 0.05 based on the Holms procedure, and an arbitrary threshold of P -value $< 10^{-5}$.

Residual permutation approach for simulations and for empirical P -value computation

To generate realistic simulation studies in which: (i) the data structure, including the exposure, covariates and outcome distributions; and (ii) their relationships, aside from the exposure–outcome association, are the same as in the real data, we used a residual permutation approach. We regressed each sleep exposure of interest X on the covariates C and estimated their effect α . We then obtained residuals, defined as:

$$\epsilon = X - C\hat{\alpha}.$$

To study type 1 error, we permuted these residuals at random to obtain $perm(\epsilon)$, and generated a sleep exposure unassociated with any of the RNA-seq measures by:

$$X_{null} = C\hat{\alpha} + perm(\epsilon).$$

We repeated this procedure 1000 times for evaluating type 1 error control. We generated simulated data under four power simulations in a similar approach, with the difference that we forced a specific correlation value between the simulated sleep exposure and a specific transcript. To this end, for a given transcript j measured on individuals $i = 1, \dots, n$, we computed the rank of each individual: $r_1(t_{1j}), \dots, r_n(t_{nj})$. To set a correlation ρ between the simulated X_ρ and transcript j we sampled $\rho \times n$ (rounded) indices from $1, \dots, n$, corresponding to $\rho \times n$ individuals for which we forced their ranks in the permuted residual values, now denoted by $perm(\epsilon)_\rho$, to be the same as their ranks in the transcript values (note that the transcript values are never changed). For the rest of the individuals, the permuted residuals are completely random. When multiple individuals have the same transcript counts (i.e. their ranks are tied), we randomly assign their ranks. For example, if 100 people have zero counts for a given transcript, each of these individuals will be equally likely to have the rank of 1, 2, \dots , or 100. The code for generating this residual permutation approach is provided in a dedicated GitHub repository <https://github.com/nkurniansyah/Olivia>.

Empirical P -values to account for the null distribution of P -values

We used the residual permutation approach, under the null hypothesis, to generate a null distribution of P -values and to compute empirical P -values. When the distribution of P -values under the null hypothesis is unknown, and specifically when it is not uniform, their values are not reliable for hypothesis testing. Alternative approaches compute ‘empirical P -values’ with the goal of generating an appropriate P -value distribution, i.e. in which an empirical P -value P_e satisfies $\Pr(P_e < 0.05 | H_0) = 0.05$ (see Supplementary Materials available online at <https://academic.oup.com/bib>).

For computing empirical P -values, we used a relatively small number of residual permutations (in comparison to the number of permutations used for computing permutation P -values) followed by transcriptome-wide association studies. We used the results of these transcriptome-wide tests under permutation to compute the null distribution of P -values, which was then used to compute the empirical P -values. We compared two types of empirical P -values: quantile empirical P -values, and Storey empirical P -values implemented in the qvalue R package [28]. The quantile empirical P -value approach was inspired by previously proposed procedures based on permutation [29] of phenotypes (rather than residuals). It estimates the null distribution of P -values non-parametrically, and the quantile empirical P -value is the quantile of the raw P -value in this distribution. The Storey empirical P -values uses the null distribution of the test statistics to identify whether a transcript is likely sampled from the null or a non-null distribution. Both implementations assume that the empirical null distribution is the same for all transcripts. We used 100 residual permutations to compute test statistics and P -values under the null and compared the empirical P -values to standard permutation P -values.

Resampling approach for binary exposure phenotypes

We compared the analysis of a continuous exposure to that of a dichotomized variable. Instead of a sleep measure, we used BMI, because it is known to have large impact of gene expression and is therefore a powerful phenotype for such a comparison. BMI was dichotomized to ‘obese’ if $BMI \geq 30 \text{ kg/m}^2$ and nonobese

otherwise. Because obesity is binary and, therefore, the residual permutation approach is not appropriate as proposed for continuous variables, we generated a binomial obesity variable based on BMI probability given covariates. Given a logistic model $\logit[p(\text{obesity}_i = 1)] = c_i^T \hat{\alpha}$, we estimated the covariates' association parameters $\hat{\alpha}$ and obtained estimated probabilities for obesity for each person $i = 1, \dots, n$ by $\hat{p}(\text{obesity}_i = 1) = \text{expit}(c_i^T \hat{\alpha})$. Based on these estimated outcome probabilities, we sampled random obesity statuses as binomial variables.

Association analysis with SDB phenotypes

Once final recommendations were developed, we implemented the proposed pipeline to study the association of AHI, MinO2 and AvgO2 jointly on RNA-seq in MESA. The associations were studied jointly for the three phenotypes: we used the multivariate Wald test [30] to obtain raw *P*-values and applied residual permutation on the three phenotypes jointly to obtain empirical *P*-values. We used the FGSEA R package [31] to apply gene set enrichment analysis on the results.

Results

MESA participant characteristics are provided in the [Supplementary Table S1](#). The distributions of the raw phenotypes AHI, MinO2 and AvgO2, and their residuals after regression on covariates are provided in [Figure 1](#), demonstrating the high non-normality. Simulations were performed after normalizing the RNA-seq data so that each library has the same size (prior to filtering), which we set to the median observed value (i.e. median normalization) in the raw reads, or 23 210 672. Results for some of the settings in simulation study 1 under TMM and size factor normalizations are provided in the Supplementary Materials available online at <https://academic.oup.com/bib>.

Simulation study 1: type 1 error analysis

After normalization, we applied filters to remove lowly expressed transcripts. There were 58 311 transcripts. After applying filters requiring that the (a) maximum read count is >10 and that (b) the proportion of individuals with zero counts for a transcript across the sample is not higher than 0.75 (see Supplementary Materials available online at <https://academic.oup.com/bib> for more information on filters), 23 004 transcripts were available for the simulation study. We used residual permutation to generate simulated SDB phenotypes that are not associated with the transcripts, but maintain the same correlation structure with the transcript and covariates. We generated 100 datasets with simulated SDB phenotypes, and performed analyses. Complete results showing the average number of false positive detection based on the existing packages limma, edgeR and DESeq2, as well as the three linear regression analyses described here, are provided in [Supplementary Figures S3–S5](#). These results include comparisons of raw *P*-values, the proposed quantile empirical *P*-values, and the empirical *P*-values provided in the *qvalue* R package [28] and for the three SDB phenotypes.

We found that the number of false positives vary with the exposure phenotypes, with analyses of MinO2 ([Figure 2](#)) generally resulting in more false positive detections than analyses of the AHI, with intermediate numbers for AvgO2 (see [Supplementary Figures S3–S5](#) available online at <https://academic.oup.com/bib>). [Figure 2](#) compares the average number of falsely discovered transcript associations when using simulated sleep phenotypes mimicking MinO2 using the residual permutation approach by

focusing on limma, edgeR, DESeq2 and linear regression applied on \log_2 of expression counts with SubHalfMin. For each method, type I error was determined using raw *P*-values and Storey empirical *P*-values, with significance thresholds based on BH FDR, local FDR (*qvalue* package) and Holms FWER. Empirical *P*-values usually reduced the number of false detections, with the method in the *qvalue* package being usually more conservative than the quantile-based empirical *P*-values method. Compared with linear regression-based approaches, DESeq2, edgeR and limma-voom had many false detections when using the raw *P*-values, even after applying multiple testing corrections. The three linear regression-based methods described here were quite similar, with the AddHalf approach often resulting in slightly more false detections. Based on these results, we chose to move forward for the next set of simulations with linear regression with SubHalfMin for handling of zero counts. In the Supplementary Materials, we report a study of false positive detections of this approach in smaller sample sizes (down to $N=30$). The results were similar.

Simulation study 2: power analysis

We performed simulations that mimic transcriptome-wide analysis to assess power. Based on simulations comparing power by transcript distributional characteristics (see Supplementary Materials section 'Study of filtering based on distributional characteristics of transcripts'), we only considered 19 742 transcripts for which no $>50\%$ of the sample had zero counts. We chose two transcripts, and for each of these and each of the sleep phenotypes, we performed 100 simulations in which we used the residual permutation approach to generate association between the sleep phenotype and the transcript with correlation $\rho = 0.3$. We performed transcriptome-wide association analysis using DESeq2, edgeR and linear regression with SubHalfMin transformation (limma-voom was not used, given its high rate of false positive detections in some of the settings in simulation study 1). For power, we always used empirical *P*-values (both types) and determined whether the specific transcript of interest passed the significance threshold based on FDR-adjusted [32] empirical *P*-value <0.05 . Power was defined as the proportion of the simulations in which the associations was significant, and was consistently higher for the linear regression-based approach compared with DESeq2 or edgeR. For linear regression, the quantile empirical *P*-values performed essentially the same as Storey empirical *P*-values, whereas Storey empirical *P*-values resulted in substantially higher statistical power when using DESeq2 and edgeR. We illustrate power comparisons in [Figure 3](#) using Storey empirical *P*-values. Power comparisons using quantile empirical *P*-values are provided in the [Supplementary Figure S9](#).

Proposed analysis approach

Based on the above simulation studies, we developed an analytic pipeline as depicted in [Figure 4](#): (i) the raw read count are normalized; (ii) filters are applied to remove lowly expressed transcripts and those for which the statistical power is low, as determined by simulations; (iii) AddHalfMin transformation is applied for each transcript separately, then log transformation is applied on all transcripts; (iv) association analyses are performed using linear regression to compute effect sizes and *P*-values; (v) permutations are computed 100 times on exposure residuals after regressing on covariates, to generate simulated phenotypes that maintain the data structure; (vi) each of 100

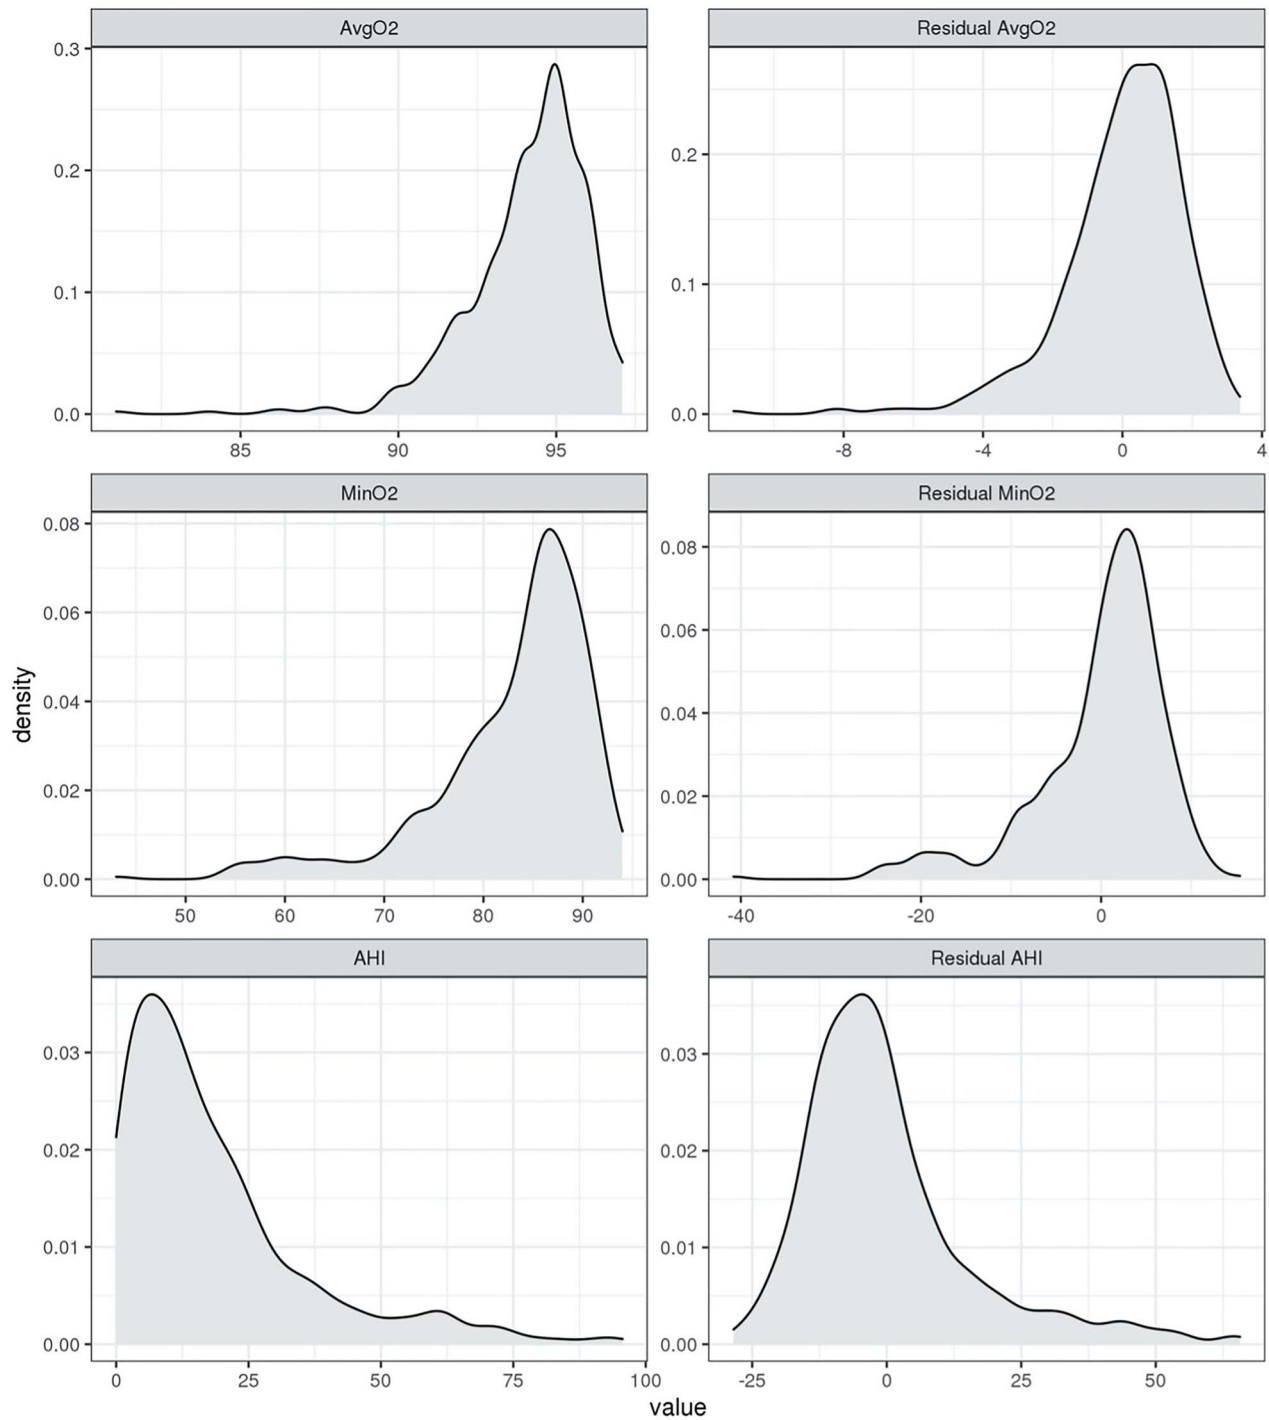


Figure 1. Distributions of the three sleep-disordered breathing exposure phenotypes used as case studies in this article. The left column provides the empirical density functions of the raw phenotypes, the right column provides the empirical density functions of their residuals after regressing on age, sex, BMI, self-reported race/ethnic group and study center.

vectors of simulated traits are analyzed using the same approach as the raw trait, generating P -values; (vii) P -values from the analysis of the 100 simulated phenotypes are combined to generate an empirical null distribution of P -values, that are used to generate empirical P -values for the raw phenotype using the `qvalue` package and (viii) multiple testing correction is applied on the empirical P -values. This pipeline is implemented in the R package `Olivia` and a corresponding Shiny app.

Comparison of analysis of continuous BMI with analysis of dichotomous obesity status

We compared the differential expression of transcripts in analysis of BMI and obesity. Residual permutation procedure was used and quantile-empirical P -values generated for both analyses. A total of 925 MESA individuals had BMI measure available and, for analysis, at least 50% nonzero transcripts were required. For obesity, several nonzero transcript thresholds were examined: 50%,




Figure 2. Average number of false positive transcript associations detected by various methods used in simulation study 1 and computed over 100 repetitions. We used the residual permutation approach to mimic the MESA dataset with the sleep phenotype MinO2. The methods reported here are linear regression (applied on log2-transformed transcript counts, with zero values replaced with SubHalfMin); DESeq2, edgeR and limmavoom. The left column provides results when using raw P-values, the middle corresponds to use of quantile empirical P-values and the right corresponds to Storey empirical P-values. We report false positive detections as those with BH FDR-adjusted P-value <0.05, Local FDR <0.05 (qvalue package) and with Holms FWER adjusted P-values <0.05. Error bars reflect the mean \pm standard error. In [Supplementary Figures S3–S5](#), we provide complete results, including for additional sleep phenotypes: AHI and AvgO2.

40% and 30%. The results were similar for all thresholds, resulting in many more identified transcript associations (446 versus 251) with continuous BMI compared with using a dichotomous trait ([Supplementary Figure S10](#)).

Computing time comparison

The computing time for transcriptome-wide association study was obtained for analyses using DESeq2, edgeR and our linear regression implementation. Using our linear regression

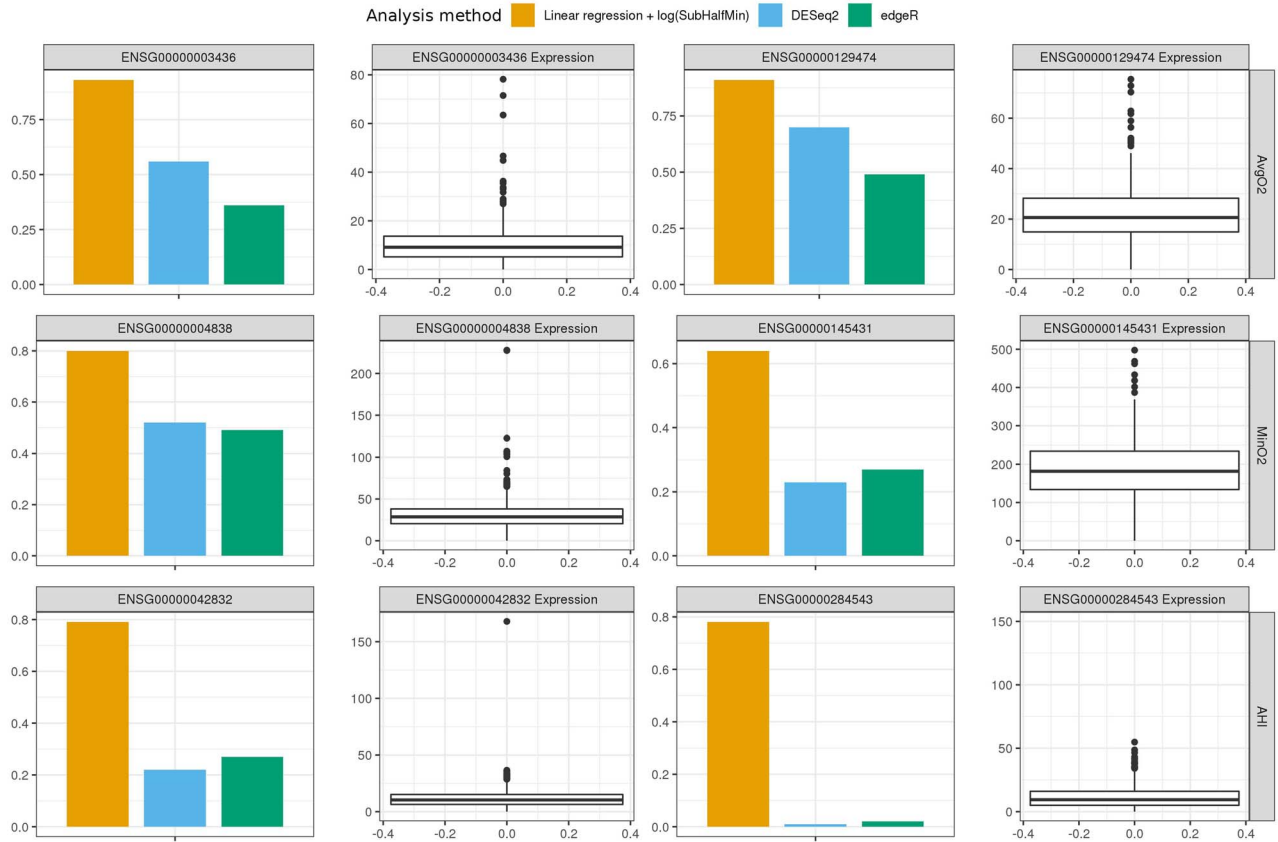


Figure 3. Estimated power for detecting a transcript simulated as associated with the three sleep traits when using Storey empirical P-values, and association is determined significant if its BH FDR-adjusted P-value is <0.05 . The transcripts were randomly selected out of available transcripts (after filtering of transcripts with 50% or higher zero counts across the sample). We compared linear regression, DESeq2, and edgeR in transcriptome-wide association analysis for each of the sleep phenotypes. For each transcript used in simulations, we show both power and the box plot of its distribution in the sample after Median normalization.

implementation on a single core, a single transcriptome-wide association study applied on ~ 19 K transcripts and $N=462$ individuals took less than a minute; when 100 transcriptome-wide association studies applied to residual permutations were included to compute empirical P-values, the time reached 7 min, and the maximum memory used was 1.3GB. In comparison, DESeq2 took 53.5 min and edgeR took 18.8 min for a single transcriptome-wide association study. The maximum memory used for DESeq2 and edgeR was similar at 3.1GB.

Association analysis with SDB phenotypes

We used the proposed pipeline to study the association of SDB phenotypes with RNA-seq. Results are provided in [Supplementary Tables S2 and S3](#) available online at <https://academic.oup.com/bib>. In brief, none of the genes passed the FDR P-value <0.05 level. The top two associations were with platelet derived growth factor C (PDGFC) gene and with Kruppel like factor 11 (KLF11) gene (both had FDR empirical P-value = 0.07). In gene set enrichment analysis, 21 Hallmark pathways [33] were significantly enriched (FDR P-value <0.05). Reassuringly, the strongest association was with the hypoxia pathway (FDR P-value = 6.12×10^{-5}). Other top pathways were of inflammatory response and heme metabolism, in agreement with past work in MESA [34].

The Olivia R package and shiny app

We developed the R package Olivia implementing the proposed procedure. Olivia is also implemented as an R Shiny app [35],

which further uses the FGSEA R package [31] to perform gene set enrichment analysis based on the results. The package and instructions for activating and using the shiny app are available in the GitHub repository <https://github.com/nkurniansyah/Olivia>. The package also provides tests of multiple exposure variables at the same time, which applies the multivariate-Wald test, and an efficient implementation of a permutation test when considering a single transcript, rather than a transcriptome-wide analysis. Finally, the repository includes code used for simulations.

Discussion

We systematically assessed the approaches for studying the association of gene expression, estimated using RNA sequencing, with continuous and non-normally distributed exposure phenotypes. We found that linear regression-based analysis performs well for continuous phenotype associations, and is computationally highly efficient. We used a residual permutation approach to study the distribution of P-values under the null of no association between the phenotypes and RNA-seq, and used this approach to further study power, and to compute empirical P-values. Notably, the residual permutation approach allows for the dataset to have the same correlation structures and associations between the phenotypes and the transcripts and covariates, while eliminating the transcript-phenotype associations. We implemented this approach in an R package and developed an R shiny app, to make our pipeline easily accessible to the research community.

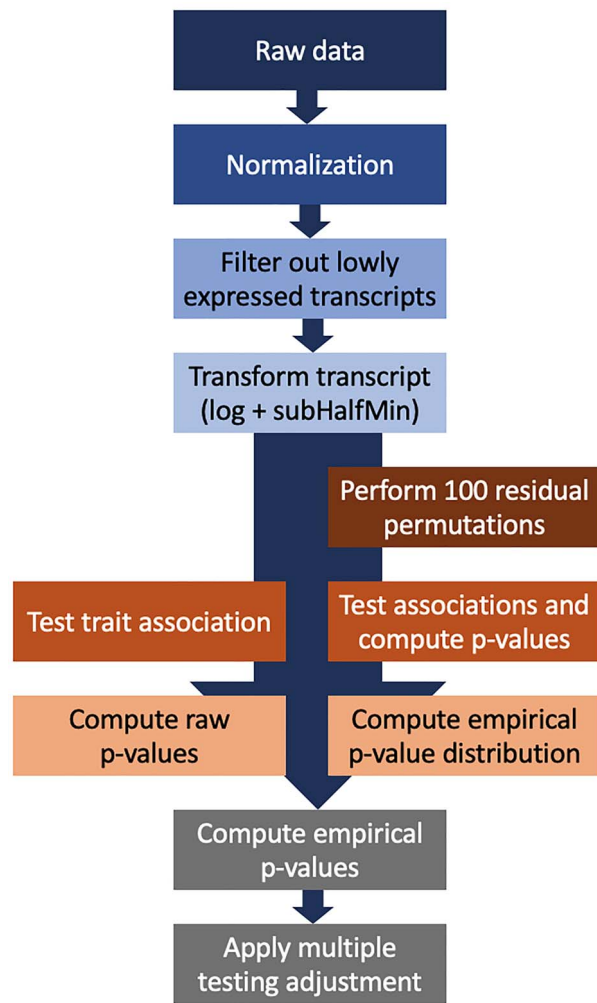


Figure 4. Analysis pipeline for association transcriptome-wide association analysis of continuous exposure phenotypes. The raw data are normalized using library-size normalization, followed by filtering of transcripts, transformation of transcript expression values, then single-transcript testing to obtain raw P -values. In parallel, residual permutation is applied under the null 100 times, and P -values are used to construct an empirical P -value distribution under the null, and to compute empirical P -values. Finally, the quantile empirical P -values are corrected for multiple testing.

Recently, van Rooij, Mandaviya [36] also performed a benchmarking study comparing analysis approaches for transcriptome-wide analysis of RNA-seq in population-based studies, including when using continuous phenotypes in association testing. Although we used similar statistical methods to theirs, we took a different analytical approach. Van Rooij *et al.* used multiple datasets to apply association analysis between a phenotype and transcripts, and assessed replication between analyses. We, on the other hand, leveraged simulations to generate data under a known association structure. In addition, we were motivated by a specific problem: highly non-normal sleep exposure measures, often leading to suboptimal control of type 1 error. Thus, it was critical to assess control of false discovery under the null hypothesis of no association between the phenotypes and RNA-seq. Notably, sleep phenotypes are less often available and there are no other large observational studies datasets to our knowledge with both RNA-seq measures and similar SDB phenotypes. Some of our findings are similar to those of van Rooij *et al.*: they also recommend using linear

regression analysis, and they also found that using a continuous phenotype is generally more powerful than dichotomizing it (in agreement with what is known from statistical literature). Similarly, they found that normalization method had very little effect on the results. However, they recommend testing all genes, while we recommend filtering transcripts with at least 50% zero counts, based on our power simulations. Additional future work is needed to evaluate various filtering criteria with relation to the phenotype distribution, and to develop methods that allow for flexible, nonlinear modeling of the association between phenotype and gene expression while remaining computationally efficient to allow for permutation analysis.

We propose to compute P -values under the null hypothesis of no association between the transcript and the exposure phenotype by permuting residuals of the exposure phenotype after regressing on covariates, and restructuring the exposure by summing the permuted residuals with the estimated mean, and thus maintain the overall data structure except for the exposure-outcome association of interest. Outside the gene expression literature, others have proposed to permute residuals rather than the outcome. For example, previous permutation methods proposed to permute residuals of the outcome after regressing on covariates [37], or to permute the residuals of the exposure phenotypes without constructing a new exposure phenotype by summing the permuted residuals with the estimated mean [38]. It will be interesting to perform a more comprehensive study of statistical permutation approaches for RNA-seq association analyses, as well as studying them in the context of mixed models.

We recommend using empirical P -values, which require 100 residual permutations, and therefore, performing 101 transcriptome-wide association analyses instead of one. Considering Supplementary Figures S3–S5 available online at <https://academic.oup.com/bib>, one can see that in most settings, linear regression methods do not have many false positive detections even when raw P -values are used. However, we chose to be more conservative by strongly protecting the analysis from false positive detections. Importantly, the linear regression analysis with empirical P -values had higher power than the other common approaches (DESeq2 and edgeR), indicating simultaneous improvement in controlling false positives and increasing power. Unfortunately, we cannot effectively estimate the FDR in these simulations. FDR is defined as the proportion of false discoveries out of all discovered (significant) associations. In simulation study 1, none of the transcripts were associated with the outcomes, so that any estimated FDR would be 100%. Under the alternative, one can suggest to use the number of wrongly discovered associations to estimate the FDR. However, many transcripts are highly correlated with the one simulated to be associated with the exposures, and are therefore associated with the exposure by design, and thus the number of transcripts falsely detected as associated with the exposure cannot be easily determined.

The empirical P -values procedure uses P -values from the entire tested transcriptome to compute the empirical null distribution. This encapsulates the assumption that the null distribution of P -values is the same for all transcripts, which is generally a limitation, but has been shown to be often acceptable since it will lead to less power, rather than increasing the number of false detections [39, 40]. An approach that does not require this assumption estimates the null distribution for P -value for each transcript separately, which is a standard permutation approach. We investigated this issue by comparing the quantile empirical P -values with the permutation P -values that use 100 000 residual permutations to estimate the null

distribution of the *P*-value of each transcript separately (see [Supplementary Figure S2](#) available online at <https://academic.oup.com/bib>). The two *P*-value distributions are very similar. Therefore, a computationally expensive permutation approach, as well as other approaches proposed by investigators, such as estimating null distributions across sets of transcripts with similar properties [39, 41], are likely unnecessary and not superior to the computationally efficient empirical *P*-values method. Another approach for estimating the null distribution of *P*-values uses the primary results, without any permutation [42, 43]. These approaches also use the assumption that the null *P*-value distribution is the same across transcripts (i.e. a shared null distribution exists). Given the computationally fast implementation of the transcriptome-wide association study, we believe that using residual permutation is beneficial because it allows for a more precise quantification of the null *P*-value distribution.

Batch effects are important to account for in studies of RNA-seq. Here, we did not study their effect on the performance of association analysis methods because it was beyond the scope of our investigation. van Rooij, Mandaviya [36] in their benchmarking study focusing on replication across cohorts, compared a few approaches for adjusting for technical covariates, including estimating and adjusting for latent confounders [44]. They concluded that inclusion of more technical adjusting covariates, including hidden confounders, increases the rate of replication between studies.

To summarize, we highlighted the problem of high false positive findings in RNA-seq data when studying the association of continuous exposure phenotypes that are highly non-normal. We developed a computationally efficient pipeline to address the false positive detection problem, and studied strategies to optimize statistical power. Our approach will be particularly useful for epidemiological studies with RNA-seq data that were not designed as disease-focused case-control studies.

Key Points

- Association analysis between a highly non-normal continuous exposure variable and RNA-seq may result in many false positive associations, whereas many *P*-values are smaller than they should be.
- Standard permutation procedures to recompute *P*-values are not appropriate for observational studies because the correlation structure between the covariates and the permuted variable are 'broken' following permutation.
- We compared approaches for association analysis combined with a 'residual permutation' scheme (permuting phenotype residuals obtained after regression on covariates) to recompute *P*-values.
- Regression-based analysis is highly computationally efficient and powerful for RNA-seq association analysis with continuous exposure variables.
- We provide recommendations and an R package *Olivia* for implementing the recommended analysis procedure.

Availability

MESA data are available through application to dbGaP. Phenotypes are available in MESA study accession phs000209.v13.p3, and RNA-seq data has been deposited and will

become available through the TOPMed-MESA study accession phs001416.v2.p1.

Supplementary data

[Supplementary data](#) are available online at *Briefings in Bioinformatics*.

Authors' contributions

TS conceptualized and drafted the manuscript and supervised the analysis. NZ performed all statistical analysis and data visualization and developed the R package and R shiny app. DN and JS performed RNA sequencing. FA and KA generated the MESA processed the sequenced RNA to generate the RNA-seq dataset. PD, YL, SSR, JIR and KDT designed the RNA-seq study in MESA. SR designed and supervised the MESA sleep ancillary study. All authors critically reviewed and approved the manuscript.

Acknowledgements

We gratefully acknowledge the studies and participants who provided biological samples and data for TOPMed.

Funding

This work was supported by the National Heart Lung and Blood Institute (grant R35HL135818). MESA and the MESA SHARE projects are conducted and supported by the National Heart, Lung and Blood Institute (NHLBI) in collaboration with MESA investigators. Support for MESA is provided by contracts 75N92020D00001, HHSN268201500003I, N01-HC-95159, 75N92020D00005, N01-HC-95160, 75N92020D00002, N01-HC-95161, 75N92020D00003, N01-HC-95162, 75N92020D00006, N01-HC-95163, 75N92020D00004, N01-HC-95164, 75N92020D00007, N01-HC-95165, N01-HC-95166, N01-HC-95167, N01-HC-95168, N01-HC-95169, UL1-TR-000040, UL1-TR-001079 and UL1-TR-001420. Also supported in part by the National Center for Advancing Translational Sciences, CTSI (grant UL1TR001881) and the National Institute of Diabetes and Digestive and Kidney Disease Diabetes Research Center (DRC; grant DK063491 to the Southern California Diabetes Endocrinology Research Center). Molecular data for the Trans-Omics in Precision Medicine (TOPMed) program was supported by the National Heart, Lung and Blood Institute (NHLBI). RNA-Seq for 'NHLBI TOPMed: Multi-Ethnic Study of Atherosclerosis (MESA)' (phs001416.v1.p1) was performed at the Northwest Genomics Center (HHSN268201600032I). Core support including centralized genomic read mapping and genotype calling, along with variant quality metrics and filtering were provided by the TOPMed Informatics Research Center (3R01HL-117626-02S1; contract HHSN268201800002I). Core support including phenotype harmonization, data management, sample-identity QC and general program coordination were provided by the TOPMed Data Coordinating Center (R01HL-120393; U01HL-120393; contract HHSN268201800001I).

References

1. Zhai W, Yao XD, Xu YF, et al. Transcriptome profiling of prostate tumor and matched normal samples by RNA-Seq. *Eur Rev Med Pharmacol Sci* 2014;**18**(9):1354–60.
2. Peng L, Bian XW, Li DK, et al. Large-scale RNA-Seq transcriptome analysis of 4043 cancers and 548 normal tissue controls across 12 TCGA cancer types. *Sci Rep* 2015;**5**:13413.
3. Kim WJ, Lim JH, Lee JS, et al. Comprehensive analysis of transcriptome sequencing data in the lung tissues of COPD subjects. *Int J Genomics* 2015;**2015**:206937.
4. Klambauer G, Unterthiner T, Hochreiter S. DEXUS: identifying differential expression in RNA-Seq studies with unknown conditions. *Nucleic Acids Res* 2013;**41**(21):e198.
5. Auer PL, Doerge RW. Statistical design and analysis of RNA sequencing data. *Genetics* 2010;**185**(2):405–16.
6. Mortazavi A, Williams BA, McCue K, et al. Mapping and quantifying mammalian transcriptomes by RNA-Seq. *Nat Methods* 2008;**5**(7):621–8.
7. Law CW, Alhamdoosh M, Su S, et al. RNA-seq analysis is easy as 1-2-3 with limma, Glimma and edge R. *F1000Research* 2016;**5**:1408.
8. Liu R, Holik AZ, Su S, et al. Why weight? Modelling sample and observational level variability improves power in RNA-seq analyses. *Nucl Acids Res* 2015;**43**(15):e97.
9. Love MI, Huber W, Anders S. Moderated estimation of fold change and dispersion for RNA-seq data with DESeq2. *Genome Biol* 2014;**15**(12):550.
10. Pimentel H, Bray NL, Puente S, et al. Differential analysis of RNA-seq incorporating quantification uncertainty. *Nat Methods* 2017;**14**(7):687–90.
11. Wolf JBW. Principles of transcriptome analysis and gene expression quantification: an RNA-seq tutorial. *Mol Ecol Resour* 2013;**13**(4):559–72.
12. Kathleen Kerr M, Churchill GA. Statistical design and the analysis of gene expression microarray data. *Genet Res* 2001;**77**(2):123–8.
13. Durbin BP, Hardin JS, Hawkins DM, et al. A variance-stabilizing transformation for gene-expression microarray data. *Bioinformatics* 2002;**18**(suppl_1):S105–10.
14. Mostafavi S, Battle A, Zhu X, et al. Normalizing RNA-sequencing data by modeling hidden covariates with prior knowledge. *PLoS One* 2013;**8**(7):e68141.
15. Conesa A, Madrigal P, Tarazona S, et al. A survey of best practices for RNA-seq data analysis. *Genome Biol* 2016;**17**:13.
16. Costa-Silva J, Domingues D, Lopes FM. RNA-Seq differential expression analysis: an extended review and a software tool. *PLoS One* 2017;**12**(12):e0190152.
17. Ge SX, Son EW, Yao R. iDEP: an integrated web application for differential expression and pathway analysis of RNA-Seq data. *BMC Bioinformatics* 2018;**19**(1):534.
18. Hrdlickova R, Toloue M, Tian B. RNA-Seq methods for transcriptome analysis. *Wiley Interdiscip Rev: RNA* 2017;**8**(1):e1364.
19. Li WV, Li JJ. Modeling and analysis of RNA-seq data: a review from a statistical perspective. *Quant Biol* 2018;**6**(3):195–209.
20. Lähnemann D, Köster J, Szczurek E, et al. Eleven grand challenges in single-cell data science. *Genome Biol* 2020;**21**(1):31.
21. Robinson MD, McCarthy DJ, Smyth GK. edgeR: a bioconductor package for differential expression analysis of digital gene expression data. *Bioinformatics* 2009;**26**(1):139–40.
22. Ritchie ME, Phipson B, Wu D, et al. Limma powers differential expression analyses for RNA-sequencing and microarray studies. *Nucl Acids Res* 2015;**43**(7):e47.
23. Dillies M-A, Rau A, Aubert J, et al. A comprehensive evaluation of normalization methods for illumina high-throughput RNA sequencing data analysis. *Brief Bioinform* 2012;**14**(6):671–83.
24. Robinson MD, Oshlack A. A scaling normalization method for differential expression analysis of RNA-seq data. *Genome Biol* 2010;**11**(3):R25.
25. Anders S, Huber W. Different expression analysis for sequence count data. *Genome Biol* 2010;**11**:R106.
26. Bild DE, Bluemke DA, Burke GL, et al. Multi-ethnic study of atherosclerosis: objectives and design. *Am J Epidemiol* 2002;**156**(9):871–81.
27. Chen X, Wang R, Zee P, et al. Racial/ethnic differences in sleep disturbances: the multi-ethnic study of atherosclerosis (MESA). *Sleep* 2015;**38**(6):877–88.
28. Storey J, Bass A, Dabney A, et al. *qvalue: Q-Value Estimation for False Discovery Rate Control*, in *R package version 2.18.0*, 2019.
29. van der Laan MJ, Hubbard AE. Quantile-function based null distribution in resampling based multiple testing. *Stat Appl Genet Mol Biol* 2006;**5**(1):Article 14.
30. Kleinbaum, D.G., L.L. Kupper, A. Nizam, et al., *Applied Regression Analysis and Other Multivariable Methods*. 5th edn 2013, Boston, MA: Cengage Learning. xix, 1051 pages.
31. Korotkevich G, Sukhov V, Sergushichev A. Fast gene set enrichment analysis. *bioRxiv* 2019;060012.
32. Benjamini Y, Hochberg Y. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J R Stat Soc: Ser B* 1995;**57**(1):289–300.
33. Liberzon A, Birger C, Thorvaldsdóttir H, et al. The molecular signatures database Hallmark gene set collection. *Cell Syst* 2015;**1**(6):417–25.
34. Sofer T, Li R, Joehanes R, et al. Low oxygen saturation during sleep reduces CD1D and RAB20 expressions that are reversed by CPAP therapy. *EBioMed* 2020;**56**:102803.
35. Chang W, Cheng J, Allaire J, et al. Shiny: Web Application Framework for R. *R package version 1.6.0*, 2021.
36. van Rooij J, Mandaviya PR, Claringbould A, et al. Evaluation of commonly used analysis strategies for epigenome- and transcriptome-wide association studies through replication of large-scale population studies. *Genome Biol* 2019;**20**(1):235.
37. Anderson MJ, Legendre P. An empirical comparison of permutation methods for tests of partial regression coefficients in a linear model. *J Stat Comput Simul* 1999;**62**(3):271–303.
38. Werft W, Benner A. Glimperm: a permutation of regressor residuals test for inference in generalized linear models. *R J* 2010;**2**(1):39.
39. Yang H, Churchill G. Estimating p-values in small microarray experiments. *Bioinformatics* 2006;**23**(1):38–43.
40. Storey, J.D. and R. Tibshirani, *SAM thresholding and false discovery rates for detecting differential gene expression in DNA microarrays*, in *The Analysis of Gene Expression Data: Methods and Software*, G. Parmigiani, et al., Editors. 2003, Springer New York: New York, NY. p. 272–90.
41. Fan J, Chen Y, Chan HM, et al. Removing intensity effects and identifying significant genes for Affymetrix arrays in macrophage migration inhibitory factor-suppressed neuroblastoma cells. *Proc Natl Acad Sci U S A* 2005;**102**(49):17751.
42. van Iterson M, van Zwet EW, Heijmans BT, et al. Controlling bias and inflation in epigenome- and transcriptome-wide association studies using the empirical null distribution. *Genome Biol* 2017;**18**(1):19.
43. Efron B. Large-scale simultaneous hypothesis testing. *J Am Stat Assoc* 2004;**99**(465):96–104.
44. Wang J, Zhao Q, Hastie T, et al. Confounder adjustment in multiple hypothesis testing. *Ann Stat* 2017;**45**(5):1863–94.