HHS Public Access

Author manuscript

IEEE/ACM Trans Comput Biol Bioinform. Author manuscript; available in PMC 2019 August 08.

Published in final edited form as:

IEEE/ACM Trans Comput Biol Bioinform. 2019 ; 16(2): 442–454. doi:10.1109/TCBB.2018.2790918.

A Unified Model for Joint Normalization and Differential Gene Expression Detection in RNA-Seq data

Kefei Liu,

Department of Radiology and Imaging Sciences, Indiana University School of Medicine, Indianapolis, IN 46202.

Jieping Ye [Senior Member, IEEE],

Department of Computational Medicine and Bioinformatics, University of Michigan, MI 48109.

Yang Yang,

School of Computer Science and Engineering, Beihang University, Beijing 100191, China.

Li Shen [Senior Member, IEEE],

Department of Radiology and Imaging Sciences, Indiana University School of Medicine, Indianapolis, IN 46202.

Hui Jiang

Department of Biostatistics, University of Michigan, MI 48109.

Abstract

The RNA-sequencing (RNA-seq) is becoming increasingly popular for quantifying gene expression levels. Since the RNA-seq measurements are relative in nature, between-sample normalization of counts is an essential step in differential expression (DE) analysis. The normalization of existing DE detection algorithms is ad hoc and performed once for all prior to DE detection, which may be suboptimal since ideally normalization should be based on non-DE genes only and thus coupled with DE detection. We propose a unified statistical model for joint normalization and DE detection of log-transformed RNA-seq data. Sample-specific normalization factors are modeled as unknown parameters in the gene-wise linear models and jointly estimated with the regression coefficients. By imposing sparsity-inducing L1 penalty (or mixed L1/L2 penalty for multiple treatment conditions) on the regression coefficients, we formulate the problem as a penalized least-squares regression problem and apply the augmented lagrangian method to solve it. Simulation studies show that the proposed model and algorithms perform better than or comparably to existing methods in terms of detection power and false-positive rate. The performance gain increases with increasingly larger sample size or higher signal to noise ratio, and is more significant when a large proportion of genes are differentially expressed in an asymmetric manner.

Index Terms—

RNA-Seq; differential expression analysis; normalization; linear regression; L1-Norm regularization; augmented Lagrangian method

1 INTRODUCTION

Ultra high-throughput sequencing of transcriptomes (RNA-seq) is a widely used method for quantifying gene expression levels due to its low cost, high accuracy and wide dynamic range for detection [1]. As of today, modern ultra high-throughput sequencing platforms can generate hundreds of millions of sequencing reads from each biological sample in a single day. RNA-seq also facilitates the detection of novel transcripts [2] and the quantification of transcripts on isoform level [3], [4]. For these reasons, RNA-seq has become the method of choice for assaying transcriptomes [5].

One major limitation of RNA-seq is that it only provides relative measurements of transcript abundances due to difference in library size (i.e., sequencing depth) between samples [6]. Normalization of RNA-seq read counts is required in gene differential expression analysis to correct for such variation between samples. A popular form of between-sample normalization is achieved by scaling raw read counts in each sample by a sample-specific factor related to library size [6], [7]. This include CPM/RPM (counts/reads per million) [8], quantile normalization [9],

[10], upper-quartile normalization [11], trimmed mean of M values [8] and DESeq normalization [12]. Also, commonly used gene expression measures, e.g., TPM (transcript per million) [13], and RPKM/FPKM (reads/fragments per kilo-base of exon per million mapped reads) [1], [2], also correct for difference in gene length within a sample [14] (the so-called within-sample normalization). In particular, the CPM/RPM (counts/reads per million) [8], TPM (transcript per million) [13], and RPKM/FPKM (reads/fragments per kilobase of exon per million mapped reads) [1], [2] for the i -th gene from the j -th sample are respectively defined as

$$\begin{aligned} \text{cpm}_{ij} &= 10^6 \frac{c_{ij}}{N_j} \\ \text{fpkm}_{ij} &= 10^9 \frac{c_{ij}}{\ell_i \cdot N_j} \\ \text{tpm}_{ij} &= 10^6 \frac{c_{ij} / \ell_i}{\sum_i c_{ij} / \ell_i} \end{aligned} \quad (1)$$

where c_{ij} is the observed read count for gene i from the j -th sample, $N_j = \sum_i c_{ij}$ is the sequencing depth in the j -th sample, and ℓ_i be the length of gene i . In this work we focus on between-sample normalization.

In traditional count-based RNA-seq analysis methods, the read counts for each gene are assumed to follow a Poisson [15] or negative binomial (NB) distribution. One issue with the

count-based RNA-seq analysis methods is that their procedures are complicated and contain many ad hoc heuristics. Moreover, the Poisson or NB distributions of counts are mathematically less tractable than the normal distribution [16], [17]. This makes count-based methods difficult to generalize to new data. Moreover, commonly used statistical methods for microarray data analysis, e.g., quality weighting of RNA samples, addition of random noise to generate technical replicates, and gene set test [16] have been designed for normally distributed data and it is unclear whether we can adapt them to count data. Also the presence of outliers is an issue that is not well addressed (addressed in a very ad hoc manner) by existing methods. To handle that, the authors of [16] take the logarithm of the raw count of reads and apply normal distribution-based statistical methods to analyze them. Note that by logarithmic transformation, the dynamic range of the RNA-seq counts is compressed such that the outlier counts are largely transformed into “normal” data. As a result, sophisticated way to detect and discard outliers [18], [19], [20] is not required.

In this paper, like in [16], [17] we work with logtransformed gene expression values and propose a unified statistical model for differential gene expression. Different from [16], [17], we model sample-specific scaling factors for between-sample normalization as unknown parameters and incorporate them into the gene-wise linear models. By imposing the sparsity-inducing penalty (L1 penalty for single treatment factor and mixed L1/L2 penalty for multiple treatment factors) on the regression coefficients and carefully choosing the tuning parameter, the model is able to achieve joint accurate detection of DE genes and between-sample normalization. To fit the model, we first eliminate sample-specific parameters using optimization argumentation to formulate the problem as a penalized linear regression problem, and then solve it with the alternating direction method of multipliers algorithm (ADMM), which is known for its fast convergence to modest accuracy [21]. Regarding the choice of tuning parameter, we theoretically derive the smallest tuning parameter $\alpha_{\max}$ that leads to all-zero solution, and thereby find a proper tuning parameter within $[0, \alpha_{\max}]$.

Note that our work is preceded by [22] which address the differential expression problem in a similar way. The difference is that the model of [22] considers only categorical or qualitative predictor/explanatory variables (treatment conditions). For example, label “0” is assigned to samples from the control group and label “1” to samples from the treatment group. While in our model, the predictor/explanatory variables can take arbitrary numeric values, and is thus a generalization of [22] from discrete to continuous predictorvariable model case. Note that the algorithm in [22] does not apply to the current numeric variable model at hand, because (i) applicability: it requires that multiple samples are present in each group but in the continuous-predictor model the concept of “group” no longer exists, or more precisely, each group is formed by only one sample; (ii) algorithmic complexity: it requires an p -dimensional exhaustive search, where p is the number of treatment conditions. When $p > 1$ (see Section 4), the algorithm is computationally very expensive.

The remainder of the paper is organized as follows. In Section 2, we formulate the problem in the context of a single treatment factor. In Section 3, we formulate the problem as a penalized simple regression problem and derive efficient ADMM algorithm to solve it, together with the estimation of noise variance and tuning parameter. In Section 4, we extend

the simple regression model to multiple linear regression model. Comparison with existing methods is presented in Section 5, followed by discussions in Section 6.

2 DATA MODEL AND PROBLEM FORMULATION

Throughout the paper, the subscript and superscript are used to index the vectors for rows and columns of a matrix, respectively. For example, the i -th row and j -th column vector of a matrix $\mathbf{A}$ is denoted as $\mathbf{a}_i$ and $\mathbf{a}^j$, respectively. Note that this does not conform to conventional notations where the subscript is used to index the columns of a matrix and the superscript is to index the rows.

2.1 Data model

Suppose there are a total of m genes measured in n samples. Let y_{ij} , $i = 1, 2, \dots, m$, $j = 1, 2, \dots, n$, be the log-transformed gene expression measurements (a small positive number is usually added before taking logarithm) for the i -th gene from the j -th sample. The following statistical model is assumed

$$y_{ij} = \beta_{i0} + \beta_i x_j + d_j + \varepsilon_{ij}, \quad (2)$$

where β_{i0} is the y -intercept for gene i , x_j , $j = 1, 2, \dots, n$, is the predictor variable that represents the treatment condition (e.g., drug dosage) for sample j , β_i is the slope or regression coefficient representing log-fold-change of expression levels of gene i with unit change of x_j , d_j is the scaling factor (e.g., log (sequencing depth) or log (library size)) for sample j for between-sample normalization [6], and $\varepsilon_{ij} \sim \mathcal{N}(0, \sigma_i^2)$ models the measurement noise. We assume that the error terms ε_{ij} are uncorrelated with the predictor variable and uncorrelated with each other (across both gene i and sample j).

In (2), we consider a single treatment condition. Extension to models with multiple treatment conditions will be discussed in Section 4.

Our main interest is to detect differentially expressed (DE) genes, i.e., whether β_i is equal to zero. If $\beta_i \neq 0$ gene i is differentially expressed across the n samples; otherwise it is not.

Remark 2.1. Since β_{i0} and d_j in (2) respectively model gene-specific factor (e.g., gene length) and sample-specific factor, model (2) is able to work with any log-transformed gene expression measures in the form of

$$y_{ij} = \log \frac{c_{ij}}{\ell_i \cdot q_j}, \quad (3)$$

where c_{ij} is the raw counts, ℓ_i is the length of gene i and q_j is the normalization factor of the j -th sample, since ℓ_i and q_j can be absorbed into β_{i0} and d_j , respectively. Note that gene expression measures of form $c_{ij}/(\ell_i q_j)$ include the raw counts (with $\ell_i = q_j = 1$), measures based

on between sample normalization only ($\ell=1$) [6], and FPKM and TPM which are shown in (1) and involve both between and within-sample normalization.

2.2 Penalized likelihood

The likelihood function based on the measured data is given by

$$L(\beta_0, \beta, d; y) = \prod_{i=1}^m \prod_{j=1}^n \frac{1}{\sqrt{2\pi\sigma_i^2}} \exp\left\{-\frac{(y_{ij} - \beta_{i0} - \beta_i x_j - d_j)^2}{2\sigma_i^2}\right\}, \quad (4)$$

where

$$\beta = \begin{pmatrix} \beta_1 \\ \beta_2 \\ \vdots \\ \beta_m \end{pmatrix}.$$

Assume that $\{\sigma_i^2\}_{i=1}^m$ are known, maximization of (4) is equivalent to minimizing the negative log-likelihood:

$$\ell(\beta_0, \beta, d) = \sum_{i=1}^m \sum_{j=1}^n \frac{1}{2\sigma_i^2} (y_{ij} - \beta_{i0} - \beta_i x_j - d_j)^2, \quad (5)$$

where we have the irrelevant constant.

In practice, we solve for $\{\sigma_i^2\}_{i=1}^m$ using an ad hoc approach, which will be described in Section 3.4.

We introduce a L1 penalty on the β_i 's,

$$p(\beta) = \alpha \|\beta\|_1 = \alpha \sum_{i=1}^m |\beta_i|. \quad (6)$$

It is well known that L1 the penalty favors sparse solutions (forces some coefficients to be exactly zero) [23]. This is reasonable since in practice many genes are not differentially expressed.

The objective function to be minimized is

$$f(\beta_0, \beta, d) = \sum_{i=1}^m \sum_{j=1}^n \frac{1}{2\sigma_i^2} (y_{ij} - \beta_{i0} - x_j \beta_i - d_j)^2 + \alpha \sum_{i=1}^m |\beta_i|. \quad (7)$$

3 ALGORITHM DEVELOPMENT

3.1 Formulation of (7) as Penalized Simple Linear Regression Model

It can be proved that the optimization problem in (7) is jointly convex in (β_0, β, d) . Therefore, the minimizer of (7) is the stationary point.

The derivative of $f(\beta_0, \beta, d)$ with respect to $d_j, j=1, 2, \dots, n$, is

$$\frac{\partial f}{\partial d_j} = \sum_{i=1}^m -\frac{1}{\sigma_i^2} (y_{ij} - \beta_{i0} - x_j \beta_i - d_j). \quad (8)$$

Setting to zero gives

$$d_j = \frac{1}{\sum_{i=1}^m \frac{1}{\sigma_i^2}} \sum_{i=1}^m \frac{1}{\sigma_i^2} (y_{ij} - \beta_{i0} - x_j \beta_i). \quad (9)$$

Model (2) is non-identifiable because we can simply add any constant to all the d_j 's, and subtract the same constant from all the β_{i0} 's, while having the same fit. To resolve this issue, we fix $d_1 = 0$. Therefore

$$d_j = d_j - d_1 = (\bar{y}_{\cdot j}^{(w)} - \bar{y}_{\cdot 1}^{(w)}) - (x_j - x_1) \bar{\beta}^{(w)}, \quad (10)$$

where

$$\bar{y}_{\cdot j}^{(w)} = \frac{1}{\sum_{i=1}^m \frac{1}{\sigma_i^2}} \sum_{i=1}^m \frac{1}{\sigma_i^2} y_{ij}, \text{ for } j = 1, 2, \dots, n, \quad (11)$$

$$\bar{\beta}^{(w)} = \frac{1}{\sum_{i=1}^m \frac{1}{\sigma_i^2}} \sum_{i=1}^m \frac{1}{\sigma_i^2} \beta_i. \quad (12)$$

Here, the superscript (w) indicates that the mean is a weighted mean instead of an unweighted one.

On the other hand, from

$$\frac{\partial f}{\partial \beta_{i0}} = -\frac{1}{\sigma_i^2} \sum_{j=1}^n (y_{ij} - \beta_{i0} - x_j \beta_i - d_j) = 0, \quad (13)$$

we have

$$\beta_{i0} = \frac{1}{n} \sum_{j=1}^n (y_{ij} - x_j \beta_i - d_j) = \bar{y}_{i\cdot} - \bar{x} \beta_i - \frac{1}{n} \sum_{j=1}^n d_j, \quad (14)$$

where

$$\bar{y}_{i\cdot} = \frac{1}{n} \sum_{j=1}^n y_{ij}, \text{ for } i = 1, 2, \dots, m, \quad (15)$$

$$\bar{x} = \frac{1}{n} \sum_{j=1}^n x_j. \quad (16)$$

From (10) we have

$$\frac{1}{n} \sum_{j=1}^n d_j = (\bar{y}^{(w)} - \bar{y}_{\cdot 1}^{(w)}) - (\bar{x} - x_1) \bar{\beta}^{(w)}, \quad (17)$$

where

$$\bar{y}^{(w)} = \frac{1}{\sum_{i=1}^m \frac{1}{\sigma_i^2}} \sum_{i=1}^m \left(\frac{1}{\sigma_i^2} \cdot \frac{1}{n} \sum_{j=1}^n y_{ij} \right). \quad (18)$$

Substituting (17) into (14) yields

$$\beta_{i0} = \bar{y}_{i\cdot} + \bar{y}_{\cdot 1}^{(w)} - \bar{y}^{(w)} + (\bar{x} - x_1) \bar{\beta}^{(w)} - \bar{x} \beta_i. \quad (19)$$

Without loss of generality, we make the following two assumptions:

Assumption 3.1.

$$\sum_{j=1}^n x_j = n\bar{x} = 0, \quad \sum_{j=1}^n x_j^2 = 1. \quad (20)$$

These assumptions are reasonable since in the model (2) the center and scaling factor of x_j 's can be absorbed into β_{i0} and $\beta_{i\cdot}$ respectively.

Then (19) simplifies to

$$\beta_{i0} = \bar{y}_{i\cdot} + \bar{y}_{i\cdot}^{(w)} - \bar{y}^{(w)} - x_1 \bar{\beta}^{(w)}. \quad (21)$$

The sum of (10) and (21) yields

$$\beta_{i0} + d_j = \bar{y}_{i\cdot} + \bar{y}_{i\cdot}^{(w)} - \bar{y}^{(w)} - x_j \bar{\beta}^{(w)}. \quad (22)$$

Substituting (22) into (7), the latter simplifies to

$$f(\beta) = \sum_{i=1}^m \frac{1}{2\sigma_i^2} \sum_{j=1}^n \left(\tilde{y}_{ij} - x_j \beta_i + x_j \bar{\beta}^{(w)} \right)^2 + \alpha \sum_{i=1}^m |\beta_i|, \quad (23)$$

where

$$\tilde{y}_{ij\cdot} = y_{ij} - \bar{y}_{i\cdot} - \bar{y}_{i\cdot}^{(w)} + \bar{y}^{(w)}. \quad (24)$$

It can be shown by straightforward calculation that $\{\tilde{y}_{ij}\}$

$$\sum_{i=1}^m \frac{1}{\sigma_i^2} \tilde{y}_{ij} = 0, \quad (25)$$

$$\sum_{j=1}^n \tilde{y}_{ij} = 0. \quad (26)$$

3.2 Model Fitting by ADMM

We propose to use the alternating direction method of multipliers (ADMM) [21] to solve (23). Although ADMM can be very slow to converge to high accuracy, it is often the case that ADMM converges to modest accuracy very fast (within a few tens of iterations) [21].

To apply the ADMM, the problem (23) is reformulated as

$$f(\beta) = \sum_{i=1}^m \frac{1}{2\sigma_i^2} \sum_{j=1}^n (\tilde{y}_{ij} - x_j\beta_i + x_j\delta_0)^2 + \alpha \sum_{i=1}^m |\beta_i|, \quad (27a)$$

subject to

$$\frac{1}{\sum_{i=1}^m \frac{1}{2\sigma_i^2}} \sum_{i=1}^m \frac{1}{\sigma_i^2} \beta_i = \delta_0. \quad (27b)$$

The augmented Lagrangian of (27) is (28) at the bottom of the page.

Step 1: Update β_i , $i=1, 2, \dots, m$:

The derivative of (28) with respect to β_i is

$$\begin{aligned} \frac{\partial L_\rho}{\partial \beta_i} = & \frac{1}{\sigma_i^2} \sum_{j=1}^n -x_j(\tilde{y}_{ij} - x_j\beta_i + x_j\delta_0) + \alpha \partial|\beta_i| \\ & + \frac{1}{\sum_{\ell=1}^m \frac{1}{2\sigma_\ell^2}} \frac{1}{\sigma_i^2} \lambda + \rho \frac{1}{\sum_{\ell=1}^m \frac{1}{2\sigma_\ell^2}} \frac{1}{\sigma_i^2} \left(\frac{1}{\sum_{\ell=1}^m \frac{1}{2\sigma_\ell^2}} \sum_{\ell=1}^m \frac{1}{\sigma_\ell^2} \beta_\ell - \delta_0 \right), \end{aligned} \quad (29)$$

where $\partial|\beta_i|$ is the subgradient of $|\beta_i|$ with respect to β_i and is defined as

$$\partial|\beta_i| = \begin{cases} 1, & \beta_i > 0 \\ -1, & \beta_i < 0 \\ [-1, 1], & \beta_i = 0 \end{cases}$$

Setting (29) equal to zero gives (30) at the bottom of the page, where T is the soft-thresholding operator:

$$T_{\sigma_i^2 \alpha} [x] = \text{sign}(x) \left(|x| - \sigma_i^2 \alpha \right)_+ = \begin{cases} x - \sigma_i^2 \alpha, & x > \sigma_i^2 \alpha \\ x + \sigma_i^2 \alpha, & x < -\sigma_i^2 \alpha \\ 0, & -\sigma_i^2 \alpha \leq x \leq \sigma_i^2 \alpha \end{cases}$$

Step 2: Update δ_0 :

The derivative of (28) with respect to δ_0 is

$$\begin{aligned} \frac{\partial L_\rho}{\partial \delta_0} &= \sum_{i=1}^m \frac{1}{\sigma_i^2} \sum_{j=1}^n x_j (\tilde{y}_{ij} - x_j \beta_i + x_j \delta_0) \quad (31) \\ &\quad - \lambda + \rho \left(\delta_0 - \frac{1}{\sum_{i=1}^m \frac{1}{\sigma_i^2}} \sum_{i=1}^m \frac{1}{\sigma_i^2} \beta_i \right). \end{aligned}$$

Setting (31) equal to zero gives

$$\begin{aligned} \delta_0 &= \frac{1}{\sum_{i=1}^m \frac{1}{\sigma_i^2} + \rho} \left(\lambda - \sum_{i=1}^m \frac{1}{\sigma_i^2} \sum_{j=1}^n x_j \tilde{y}_{ij} \right) + \frac{1}{\sum_{i=1}^m \frac{1}{\sigma_i^2}} \sum_{i=1}^m \frac{1}{\sigma_i^2} \beta_i \quad (32) \\ &= \frac{1}{\sum_{i=1}^m \frac{1}{\sigma_i^2} + \rho} \lambda + \frac{1}{\sum_{i=1}^m \frac{1}{\sigma_i^2}} \sum_{i=1}^m \frac{1}{\sigma_i^2} \beta_i, \end{aligned}$$

where the second equality is due to (32) where the second equality is due to (25). Step 3: Update λ :

$$\lambda^{\text{new}} = \lambda^{\text{old}} + \rho \left(\frac{1}{\sum_{i=1}^m \frac{1}{\sigma_i^2}} \sum_{i=1}^m \frac{1}{\sigma_i^2} \beta_i - \delta_0 \right) \quad (33)$$

The model fitting algorithm is described in Algorithm 1.

3.3 Estimation of Tuning Parameter α

Eq. (23) can be expressed in matrix form as

$$f(\beta) = \frac{1}{2} \left\| \Sigma^{1/2} (\tilde{Y} - M\beta x^T) \right\|_F^2 + \alpha \|\beta\|_1, \quad (34)$$

where

$$\mathbf{\Sigma} = \text{diag}\{\boldsymbol{\sigma}\}, \quad (35)$$

with

$$\boldsymbol{\sigma} = (1/\sigma_1^2 \ 1/\sigma_2^2 \ \dots \ 1/\sigma_m^2)^T, \quad (36)$$

and

$$\mathbf{M} = \begin{pmatrix} 1 & 0 & \dots & 0 \\ 0 & 1 & \dots & 0 \\ 0 & 0 & \ddots & 0 \\ 0 & 0 & \dots & 1 \end{pmatrix} - \frac{1}{\sum_{i=1}^m \frac{1}{\sigma_i^2}} \begin{pmatrix} 1/\sigma_1^2 & 1/\sigma_2^2 & \dots & 1/\sigma_m^2 \\ 1/\sigma_1^2 & 1/\sigma_2^2 & \dots & 1/\sigma_m^2 \\ \vdots & \vdots & \ddots & \vdots \\ 1/\sigma_1^2 & 1/\sigma_2^2 & \dots & 1/\sigma_m^2 \end{pmatrix} = \mathbf{I}_m - \frac{1}{\sum_{i=1}^m \frac{1}{\sigma_i^2}} \mathbf{1}_m \mathbf{1}_m^T. \quad (37)$$

After expansion, (34) becomes

$$f(\boldsymbol{\beta}) = \frac{1}{2} \|\mathbf{\Sigma}^{1/2} \tilde{\mathbf{Y}}\|_F^2 - \boldsymbol{\beta}^T \mathbf{M}^T \mathbf{\Sigma} \tilde{\mathbf{Y}} \mathbf{x} + \frac{1}{2} \boldsymbol{\beta}^T \mathbf{M}^T \mathbf{\Sigma} \mathbf{M} \boldsymbol{\beta} + \alpha \|\boldsymbol{\beta}\|_1, \quad (38)$$

where we exploit assumption $\mathbf{x}^T \mathbf{x} = \mathbf{1}$.

Since $\frac{1}{2} \boldsymbol{\beta}^T \mathbf{M}^T \mathbf{\Sigma} \mathbf{M} \boldsymbol{\beta} \geq 0$ with equality occurring at $\boldsymbol{\beta} = \mathbf{0}$, it is shown that $\hat{\boldsymbol{\beta}} = \mathbf{0}$ is the minimizer of $f(\boldsymbol{\beta})$ when

$$\alpha \geq \|\mathbf{M}^T \mathbf{\Sigma} \tilde{\mathbf{Y}} \mathbf{x}\|_\infty = \max_{1 \leq i \leq m} |\mathbf{m}^i{}^T \mathbf{\Sigma} \tilde{\mathbf{Y}} \mathbf{x}|, \quad (39)$$

where $\mathbf{m}^i$ denotes the i -th column of $\mathbf{M}$ in (37).

$$\begin{aligned} L_\rho(\boldsymbol{\beta}, \delta_0, \lambda) &= \sum_{i=1}^m \frac{1}{2\sigma_i^2} \sum_{j=1}^n (\tilde{y}_{ij} - x_j \beta_i + x_j \delta_0)^2 + \alpha \sum_{i=1}^m |\beta_i| + \lambda \left(\frac{1}{\sum_{i=1}^m \frac{1}{\sigma_i^2}} \sum_{i=1}^m \frac{1}{\sigma_i^2} \beta_i - \delta_0 \right) \\ &\quad + \frac{\rho}{2} \left(\frac{1}{\sum_{i=1}^m \frac{1}{\sigma_i^2}} \sum_{i=1}^m \frac{1}{\sigma_i^2} \beta_i - \delta_0 \right)^2 \end{aligned} \quad (28)$$

$$\beta_i = \frac{\sigma_i^2 \left(\sum_{\ell=1}^m \sigma_{\ell}^{-2} \right)^2}{\sigma_i^2 \left(\sum_{\ell=1}^m \sigma_{\ell}^{-2} \right)^2 + \rho} T \sigma_i^2 \alpha \left[\left(\sum_{j=1}^n x_j \tilde{y}_{ij} + \delta_0 \right) - \frac{\rho}{\sum_{\ell=1}^m \frac{1}{\sigma_{\ell}^2}} \left(\frac{1}{\sum_{\ell=1}^m \frac{1}{\sigma_{\ell}^2}} \sum_{\ell \neq i} \frac{1}{\sigma_{\ell}^2} \beta_{\ell} - \delta_0 + \frac{\lambda}{\rho} \right) \right] \quad (30)$$

Algorithm 1 ADMM algorithm for fitting the simple linear regression model

Input: Log-transformed gene expression measurements: $\{ \{y_{ij}\}_{j=1}^n \}_{i=1}^m$, predictor variables: $\{x_j\}_{j=1}^n$ and estimated noise variance: $\{\sigma_i^2\}_{i=1}^m$.

1: Transform data. Normalize $\{x_j\}_{j=1}^n$ to zero mean and unit norm:

$$\tilde{x}_j = \frac{x_j - \bar{x}}{\sqrt{\sum_{j=1}^n (x_j - \bar{x})^2}}, \text{ with } \bar{x} = \frac{1}{n} \sum_{j=1}^n x_j.$$

Center y_{ij} to zero mean over row index i and column index j :

$$\tilde{y}_{ij} = y_{ij} - \bar{y}_{i\cdot} - \bar{y}_{\cdot j} + \bar{y}^{(w)},$$

where $\bar{y}_{i\cdot}$, $\bar{y}_{\cdot j}^{(w)}$ and $\bar{y}^{(w)}$ are defined in (15), (11) and (18), respectively.

2: Set the penalty parameter to $\rho = 1$ [21]; select the tuning parameter α according to Section 3.3.

3: **Initialization:** $k = 0$; randomly initialize $\beta_i = \beta_i^0$, $i = 1, 2, \dots, m$, $\delta_0 = \delta_0^0$, and $\lambda = \lambda^0$.

4: **repeat**

5: **Update** β_i , $i = 1, 2, \dots, m$:

$$\beta_i^{k+1} = \frac{\sigma_i^2 \left(\sum_{\ell=1}^m \sigma_{\ell}^{-2} \right)^2}{\sigma_i^2 \left(\sum_{\ell=1}^m \sigma_{\ell}^{-2} \right)^2 + \rho} T_{\alpha} \left[\left(\sum_{j=1}^n \tilde{x}_j \tilde{y}_{ij} + \delta_0^k \right) - \frac{\rho}{\sum_{\ell=1}^m \frac{1}{\sigma_{\ell}^2}} \left(\frac{1}{\sum_{\ell=1}^m \frac{1}{\sigma_{\ell}^2}} \sum_{\ell \neq i} \frac{1}{\sigma_{\ell}^2} \beta_{\ell}^k - \delta_0^k + \frac{\lambda^k}{\rho} \right) \right]$$

6: **Update** δ_0 :

$$\delta_0^{k+1} = \frac{1}{\sum_{i=1}^m \frac{1}{\sigma_i^2}} \lambda^k + \frac{1}{\sum_{i=1}^m \frac{1}{\sigma_i^2}} \sum_{i=1}^m \frac{1}{\sigma_i^2} \beta_i^{k+1}$$

7: **Update** λ :

$$\lambda^{k+1} = \lambda^k + \rho \left(\frac{1}{\sum_{i=1}^m \frac{1}{\sigma_i^2}} \sum_{i=1}^m \frac{1}{\sigma_i^2} \beta_i^{k+1} - \delta_0^{k+1} \right)$$

8: $k = k + 1$;

9: **until** convergence or maximum number of iterations is reached.

Output: $\beta_i = \beta_i^*$, $i = 1, 2, \dots, m$, $\beta^{(w)} = \delta_0^*$, and

$$\tilde{\beta}_{i0} = \tilde{y}_{i\cdot} + \tilde{y}_1^{(w)} - \tilde{y}^{(w)} - \tilde{x}_1 \tilde{\beta}^{(w)}, \quad i = 1, 2, \dots, m$$

$$d_1 = 0, \quad d_j = \left(\tilde{y}_{\cdot j}^{(w)} - \tilde{y}_1^{(w)} \right) - \left(\tilde{x}_j - \tilde{x}_1 \right) \tilde{\beta}^{(w)}, \quad j = 2, \dots, n$$

Recover the original parameter space:

$$\beta_{i0} = \tilde{\beta}_{i0} - \frac{\tilde{\beta}_1 \bar{x}}{\sqrt{\sum_{j=1}^n (x_j - \bar{x})^2}}, \quad \beta_j^* = \frac{\tilde{\beta}_j}{\sqrt{\sum_{j=1}^n (x_j - \bar{x})^2}}, \quad i = 1, 2, \dots, m.$$

Note that

$$M^T \Sigma \tilde{Y} = \left(I_m - \frac{1}{\sum_{i=1}^m \frac{1}{\sigma_i^2}} \sigma \mathbf{1}_m^T \right) \Sigma \tilde{Y} = \Sigma \tilde{Y}, \quad (40)$$

where the last equality holds because $\mathbf{1}_m^T \Sigma \tilde{Y} = \mathbf{0}$ due to (25)

Substituting (40) into (39) yields

$$\alpha_{\max} = \|\Sigma \tilde{Y} \mathbf{x}\|_{\infty} = \max_i \left| \frac{1}{\sigma_i^2} \mathbf{x}^T \tilde{\mathbf{y}}_i \right|. \quad (41)$$

Our strategy is to first sort $\left| \frac{1}{\sigma_1^2} \mathbf{x}^T \tilde{\mathbf{y}}_1 \right|, \left| \frac{1}{\sigma_2^2} \mathbf{x}^T \tilde{\mathbf{y}}_2 \right|, \dots, \left| \frac{1}{\sigma_N^2} \mathbf{x}^T \tilde{\mathbf{y}}_m \right|$ in ascending order, from least to greatest, and then set α as the P -th percentile ($0 < P < 100$) of the m ordered value. We set $P=5$ in Section 5.

3.4 Maximum likelihood estimation of $\{\sigma_i^2\}_{i=1}^m$

To solve for $\{\sigma_i^2\}_{i=1}^m$, consider the negative log-likelihood function in (4) with $\{\sigma_i^2\}_{i=1}^m$ being unknown parameters as well:

$$\ell(\beta_0, \beta, d, \{\sigma_i^2\}_{i=1}^m) = \sum_{i=1}^m \left[\frac{n}{2} \log(2\pi\sigma_i^2) + \frac{1}{2\sigma_i^2} \sum_{j=1}^n (y_{ij} - \beta_{i0} - x_j\beta_i - d_j)^2 \right] \quad (42)$$

Taking the partial derivatives of $\ell(\cdot)$ with respect to d_j and β_{i0} and setting the results to zero, we arrive at (10) and (21) respectively. The sum of (10) and (21) gives (22).

Taking the partial derivative of $\ell(\cdot)$ with respect to β_i and setting the result to zero, we have

$$\beta_i = \sum_{j=1}^n x_j y_{ij} - \sum_{j=1}^n x_j (\beta_{i0} + d_j). \quad (43)$$

Substituting (22) into (43) yields

$$\beta_i = \sum_{j=1}^n x_j y_{ij} - \frac{1}{\sum_{i=1}^m \frac{1}{\sigma_i^2}} \sum_{i=1}^m \frac{1}{\sigma_i^2} \sum_{j=1}^n x_j y_{ij} + \bar{\beta}^{(w)}, \quad (44)$$

where $\bar{\beta}^{(w)}$ is defined in (12).

Taking the partial derivative of σ_i^2 and setting the result to zero gives

$$\sigma_i^2 = \frac{1}{n} \sum_{j=1}^n (y_{ij} - \beta_{i0} - x_j\beta_i - d_j)^2. \quad (45)$$

Substituting (22) into (45) yields

$$\sigma_i^2 = \frac{1}{n} \sum_{j=1}^n (y_{ij} - \bar{y}_{i\cdot} - \bar{y}_{\cdot j}^{(w)} + \bar{y}^{(w)} - x_j\beta_i + x_j\bar{\beta}^{(w)})^2, \quad (46)$$

where $\bar{y}_{i\cdot}$, $\bar{y}_{\cdot j}^{(w)}$ and $\bar{y}^{(w)}$ are defined in (15), (11) and (18), respectively.

Given initial estimates for $\bar{\beta}^{(w)}$ and $\{\sigma_i^2\}_{i=1}^m$ we can alternate equations (44), (46) and (12) iteratively to gradually refine the estimates for β_i and σ_i^2 , as shown in Algorithm 2.

Algorithm 2 Maximum likelihood estimation of $\{\sigma_i^2\}_{i=1}^m$ for simple linear regression model

Input: Log-transformed gene expression measurements: $\{y_{ij}\}_{i=1}^m\}_{j=1}^n$ and predictor variables: $\{x_j\}_{j=1}^n$.

1: Normalize $\{x_j\}_{j=1}^n$ to zero mean and unit norm:

$$x_j \leftarrow \frac{x_j - \bar{x}}{\sqrt{\sum_{j=1}^n (x_j - \bar{x})^2}}, \text{ with } \bar{x} := \frac{1}{n} \sum_{j=1}^n x_j.$$

2: Initialization: $\bar{\beta}^{(w)} = 0, \sigma_1^2 = \sigma_2^2 = \dots = \sigma_m^2 = 1$.

3: **repeat**

4: Update $\beta_i, i = 1, 2, \dots, m$, according to (44);

5: Update $\sigma_i^2, i = 1, 2, \dots, m$, according to (46);

6: Update $\bar{\beta}^{(w)}$ according to (12);

7: **until** convergence or maximum number of iterations is reached.

Output: $\hat{\sigma}_i^2 = \sigma_i^2, i = 1, 2, \dots, m$.

To obtain a robust estimate for σ_i^2 , we further take the weighted average of $\hat{\sigma}_i^2$ and the estimated mean variance across all the genes. That is

$$\hat{\sigma}_i'^2 = (1 - w)\hat{\sigma}_i^2 + w\bar{\sigma}^2 \quad (47)$$

where

$$\bar{\sigma}^2 = \frac{1}{m} \sum_{i=1}^m \hat{\sigma}_i^2, \quad (48)$$

and the weight w is calculated using the following formula which is derived based on an empirical Bayes approach [24]

$$w = \frac{2(m-1)}{n+1} \left(\frac{1}{m} + \frac{(\bar{\sigma}^2)^2}{\sum_{i=1}^m (\hat{\sigma}_i^2 - \bar{\sigma}^2)^2} \right). \quad (49)$$

This kind of variance estimation approach is widely used in differential gene expression analysis with small sample sizes [25], [26]. The estimated variances $\hat{\sigma}_i'^2$ can then used in Algorithm 1 to solve for $\{\beta_i\}_{i=1}^m$. *Remark 3.1.* In the special case of $\sigma_1^2 = \sigma_2^2 = \dots = \sigma_m^2 = \sigma^2$, it no longer requires to estimate σ^2 since the unknown σ^2 in (7) can be absorbed into the tuning parameter α .

4 EXTENSION TO MULTIPLE LINEAR REGRESSION MODEL AND ALGORITHM DEVELOPMENT

In the multiple linear regression model, each response or outcome is modeled by $p > 1$ predictors:

$$y_{ij} = \beta_{i0} + \boldsymbol{\beta}_i^T \mathbf{x}_j + d_j + \varepsilon_{ij}, \quad (50)$$

where

$$\boldsymbol{\beta}_i = \begin{pmatrix} \beta_{i1} \\ \beta_{i2} \\ \vdots \\ \beta_{ip} \end{pmatrix} \in \mathbb{R}^{p \times 1} \quad (51)$$

is a vector of regression coefficients representing log-fold-change of expression levels of gene i between treatment conditions, and

$$\mathbf{x}_j = \begin{pmatrix} x_{j1} \\ x_{j2} \\ \vdots \\ x_{jp} \end{pmatrix} \in \mathbb{R}^{p \times 1} \quad (52)$$

is a vector of predictors representing the treatment conditions (drug dosage, blood pressure, age, BMI, etc.) for sample j , and β_{i0} , d_j and $\varepsilon_{ij} \sim \mathcal{N}(0, \sigma_i^2)$ are the y -intercept, scaling factor for sample j and measurement noise, respectively. We assume that the error terms ε_{ij} are uncorrelated with all the predictor variables and uncorrelated with each other.

The likelihood function based on the observed data is given by

$$L(\beta_0, \{\boldsymbol{\beta}_i\}_{i=1}^m, \mathbf{d}; \mathbf{Y}) = \prod_{i=1}^m \prod_{j=1}^n \frac{1}{\sqrt{2\pi\sigma_i^2}} \exp \left\{ -\frac{(y_{ij} - \beta_{i0} - \boldsymbol{\beta}_i^T \mathbf{x}_j - d_j)^2}{2\sigma_i^2} \right\}. \quad (53)$$

Assume that $\{\sigma_i^2\}_{i=1}^m$ are known, maximization of (53) leads to minimizing the negative log-likelihood:

$$\ell(\beta_0, \{\beta_i\}_{i=1}^m, d) = \sum_{i=1}^m \sum_{j=1}^n \frac{1}{2\sigma_i^2} (y_{ij} - \beta_{i0} - \beta_i^T \mathbf{x}_j - d_j)^2 \quad (54)$$

The objective function to be minimized is

$$f(\beta_0, \{\beta_i\}, d) = \sum_{i=1}^m \sum_{j=1}^n \frac{1}{2\sigma_i^2} (y_{ij} - \beta_{i0} - \mathbf{x}_j^T \beta_i - d_j)^2 + \sum_{i=1}^m p(\beta_i). \quad (55)$$

Below we introduce two types of penalty function $p(\beta_i)$.

1. Type I penalty:

$$p(\beta_i) = \alpha |\beta_{ip}|. \quad (56)$$

Gene i is differentially expressed if $\beta_{ip} \neq 0$ and not otherwise. This penalty is for the applications where one covariate is of main interest (e.g., treatment) while we want to adjust for all possible effects of other confounding covariates (e.g., age, gender, etc).

1. Type II penalty:

$$p(\beta_i) = \alpha \|\beta_i\|. \quad (57)$$

Gene i is differentially expressed if $\beta_i \neq 0$ and not otherwise. This penalty is for the applications where all covariates are of interest and we want to identify the genes for which at least one covariate has an effect.

It can be proved that the optimization problem (55) with penalty (56) or (57) is jointly convex in $(\beta_0, \{\beta_i\}, d)$.

Assume that

$$\sum_{j=1}^n \mathbf{x}_j = 0, \quad (58)$$

and set $d_1=0$. using similar argumentation as in Section 3.1 to eliminate β_0 and d , we simplify (55) to

$$f(\{\beta_i\}) = \sum_{i=1}^m \frac{1}{2\sigma_i^2} \sum_{j=1}^n (\tilde{y}_{ij} - \mathbf{x}_j^T \beta_i + \mathbf{x}_j^T \bar{\beta}^{(w)})^2 + \sum_{i=1}^m p(\beta_i), \quad (59)$$

where $\tilde{y}_{ij}$ is the same as that in (24), and

$$\bar{\beta}^{(w)} = \frac{1}{\sum_{i=1}^m \frac{1}{2\sigma_i^2}} \sum_{i=1}^m \frac{1}{2\sigma_i^2} \beta_i. \quad (60)$$

4.1 Regression with type I penalty: Model fitting by ADMM

To apply the ADMM, we reformulate the Type I penalized regression problem as

$$f(\{\beta_i\}, \delta_0) = \sum_{i=1}^m \frac{1}{2\sigma_i^2} \sum_{j=1}^n (\tilde{y}_{ij} - \mathbf{x}_j^T \beta_i + \mathbf{x}_j^T \delta_0)^2 + \alpha \sum_{i=1}^m |\beta_{ip}|, \quad (61a)$$

subject to

$$\frac{1}{\sum_{i=1}^m \frac{1}{2\sigma_i^2}} \sum_{i=1}^m \frac{1}{2\sigma_i^2} \beta_i = \delta_0. \quad (61b)$$

The augmented Lagrangian of (61) is (62) at the bottom of the page.

Step 1: Update β_i , $i = 1, 2, \dots, m$.

Taking the partial derivative of (62) with respect to β_i and setting the result to zero gives

$$\left[\frac{1}{\sigma_i^2} \mathbf{X}^T \mathbf{X} + \frac{\rho}{\sigma_i^4} \frac{1}{\left(\sum_{\ell=1}^m \frac{1}{\sigma_\ell^2} \right)^2} \mathbf{I}_p \right] \beta_i + \alpha \begin{pmatrix} 0 \\ \vdots \\ 0 \\ \partial |\beta_{ip}| \end{pmatrix} = \mathbf{v}_i, \quad (63)$$

where

$$X = \begin{pmatrix} \mathbf{x}_1^T \\ \mathbf{x}_2^T \\ \vdots \\ \mathbf{x}_n^T \end{pmatrix} = \begin{pmatrix} x_{11} & x_{12} & \cdots & x_{1p} \\ x_{21} & x_{22} & \cdots & x_{2p} \\ \vdots & \vdots & \ddots & \vdots \\ x_{n1} & x_{n2} & \cdots & x_{np} \end{pmatrix} \in \mathbb{R}^{n \times p}, \quad (64)$$

$|\beta_{ip}|$ is the subgradient of $|\beta_{ip}|$ with respect to β_{ip} , and

$$\begin{aligned} v_i &= \frac{1}{\sigma_i^2} \left(\sum_{j=1}^n x_j \tilde{y}_{ij} + X^T X \delta_0 \right) \\ &\quad - \frac{\rho}{\sigma_i^2} \frac{1}{\sum_{\ell=1}^m \frac{1}{\sigma_{\ell}^2}} \left(\frac{1}{\sum_{\ell=1}^m \frac{1}{\sigma_{\ell}^2}} \sum_{\ell \neq i} \frac{1}{\sigma_{\ell}^2} \beta_{\ell} - \delta_0 + \frac{\lambda}{\rho} \right). \end{aligned} \quad (65)$$

Given matrix partition in the following form:

$$\begin{aligned} \frac{1}{\sigma_i^2} X^T X + \frac{\rho}{\sigma_i^4} \frac{1}{\left(\sum_{\ell=1}^m \frac{1}{\sigma_{\ell}^2} \right)^2} \mathbf{I}_p &= \mathbf{Q} = \begin{pmatrix} \mathbf{Q}_{11} & \mathbf{q} \\ \mathbf{q}^T & q_{pp} \end{pmatrix}, \\ \beta_i &= \begin{pmatrix} \beta_i^- \\ \beta_{ip} \end{pmatrix}, v_i = \begin{pmatrix} v_i^- \\ v_{ip} \end{pmatrix}, \end{aligned}$$

where $\mathbf{Q}_{11}$ is the submatrix, of $\mathbf{Q}$, with last, row and last column deleted, from (63) we have

$$\mathbf{Q}_{11} \beta_i^- + \mathbf{q} \beta_{ip} = v_i^- \quad (66)$$

$$\mathbf{q}^T \beta_i^- + q_{pp} \beta_{ip} + \alpha \partial |\beta_{ip}| = v_{ip}. \quad (67)$$

From (66) it follows

$$\beta_i^- = \mathbf{Q}_{11}^{-1} (v_i^- - \mathbf{q} \beta_{ip}). \quad (68)$$

Substituting (68) into (67) yields

$$\beta_{ip} = \frac{1}{q_{pp} - \mathbf{q}^T \mathbf{Q}_{11}^{-1} \mathbf{q}} T_{\alpha} \left[v_{ip} - \mathbf{q}^T \mathbf{Q}_{11}^{-1} \mathbf{v}_i^- \right]. \quad (69)$$

Step 2: Update δ_0 :

Taking the derivative of (62) with respect to δ_0 and setting the result to zero gives

$$\delta_0 = \left(\sum_{i=1}^m \frac{1}{\sigma_i^2} \mathbf{X}^T \mathbf{X} + \rho \mathbf{I}_p \right)^{-1} \lambda + \frac{1}{\sum_{i=1}^m \frac{1}{\sigma_i^2}} \sum_{i=1}^m \frac{1}{\sigma_i^2} \beta_i, \quad (70)$$

where we have exploited (25).

Step 3: Update λ :

$$\lambda^{\text{new}} = \lambda^{\text{old}} + \rho \left(\frac{1}{\sum_{i=1}^m \frac{1}{\sigma_i^2}} \sum_{i=1}^m \frac{1}{\sigma_i^2} \beta_i - \delta_0 \right). \quad (71)$$

The model fitting algorithm is described in Algorithm 3.

4.2 Regression with type II penalty: Model fitting by ADMM

The Type II penalized regression problem is reformulated as

$$f(\{\beta_i\}, \delta_0) = \sum_{i=1}^m \frac{1}{2\sigma_i^2} \sum_{j=1}^n (\tilde{y}_{ij} - \mathbf{x}_j^T \beta_i + \mathbf{x}_j^T \delta_0)^2 + \alpha \sum_{i=1}^m \|\beta_i\|, \quad (72a)$$

subject to

$$\frac{1}{\sum_{i=1}^m \frac{1}{\sigma_i^2}} \sum_{i=1}^m \frac{1}{\sigma_i^2} \beta_i = \delta_0. \quad (72b)$$

The augmented Lagrangian of (72) is (73) at the bottom of the page.

Step 1: Update β_i $i = 1, 2, \dots, m$:

The relevant terms to compute the derivatives of (73) with respect to β_i is (74) at the bottom of the page, where c is an irrelevant constant which does not depend on β_i and $\mathbf{v}_i$ is defined in (65).

$$L_\rho(\{\beta_i\}, \delta_0, \lambda) = \sum_{i=1}^m \frac{1}{2\sigma_i^2} \sum_{j=1}^n (\tilde{y}_{ij} - \mathbf{x}_j^T \beta_i + \mathbf{x}_j^T \delta_0)^2 + \alpha \sum_{i=1}^m |\beta_{ip}| \quad (62)$$

$$+ \lambda^T \left(\frac{1}{\sum_{i=1}^m \frac{1}{\sigma_i^2}} \sum_{i=1}^m \frac{1}{\sigma_i^2} \beta_i - \delta_0 \right) + \frac{\rho}{2} \left\| \frac{1}{\sum_{i=1}^m \frac{1}{\sigma_i^2}} \sum_{i=1}^m \frac{1}{\sigma_i^2} \beta_i - \delta_0 \right\|^2.$$

It can be shown that when $\|\mathbf{v}_i\| = \alpha$ then $\beta_i = \mathbf{0}$; otherwise denote the eigendecomposition of $\mathbf{X}^T \mathbf{X}$ as $\mathbf{X}^T \mathbf{X} = \mathbf{U} \mathbf{D} \mathbf{U}^T$, we have that minimization of (74) is equivalent to

$$\min_{\beta_i} \frac{1}{2} \|\mathbf{Z}_i \beta_i - \mathbf{b}_i\|^2 + \alpha \|\beta_i\|, \quad (75a)$$

where

$$\mathbf{Z}_i = \left[\frac{1}{\sigma_i^2} \mathbf{D} + \frac{\rho}{\sigma_i^4} \frac{1}{\left(\sum_{\ell=1}^m \frac{1}{\sigma_\ell^2} \right)^2} \mathbf{I}_p \right]^{1/2} \mathbf{U}^T, \quad (75b)$$

$$\mathbf{b}_i = \left[\frac{1}{\sigma_i^2} \mathbf{D} + \frac{\rho}{\sigma_i^4} \frac{1}{\left(\sum_{\ell=1}^m \frac{1}{\sigma_\ell^2} \right)^2} \mathbf{I}_p \right]^{-1/2} \mathbf{U}^T \mathbf{v}_i. \quad (75c)$$

As in [27], we use a coordinate descent procedure to optimize (75). For each s , given the estimate of $\{\hat{\beta}_{i\ell}\}_{\ell \neq s}$ can be estimated by solving

$$\min_{\beta_{is}} \frac{1}{2} \|\mathbf{z}^s \beta_{is} - r_i^{(s)}\|^2 + \alpha \sqrt{\beta_{is}^2 + \sum_{\ell \neq s} \hat{\beta}_{i\ell}^2}, \quad (76)$$

where

$$\mathbf{r}_i^{(s)} = \mathbf{b}_i - \sum_{\ell \neq s} \mathbf{z}^\ell \hat{\beta}_{i\ell}. \quad (77)$$

We solve (76) via a one-dimensional search. Note that the solution to (76) falls between 0 and $\beta_{i\ell}^0 = \mathbf{z}^{sT} \mathbf{r}_i^{(s)} / \|\mathbf{z}^s\|^2$, the ordinary least-squares estimate. We can use the optimize function in the R package, or fminbnd function in MAT-LAB, which performs one-dimensional search based on golden section search and successive parabolic interpolation.

After updating $\{\beta_i\}_{i=1}^m$, the updates of $\boldsymbol{\delta}_0$ and $\boldsymbol{\lambda}$ out to be the same as that in Section 4.1.

The model fitting algorithm is described in Algorithm 4.

4.3 Estimation of Tuning Parameter α

Eq. (59) can be expressed in matrix form as

$$f(\mathbf{B}) = \frac{1}{2} \left\| \boldsymbol{\Sigma}^{1/2} (\tilde{\mathbf{Y}} - \mathbf{M}\mathbf{B}\mathbf{X}^T) \right\|_F^2 + p(\mathbf{B}), \quad (78)$$

where $\mathbf{M}$ and $\mathbf{X}$ are respectively defined in (37) and (64), and

$$\mathbf{B} = \begin{pmatrix} \beta_1^T \\ \beta_2^T \\ \vdots \\ \beta_m^T \end{pmatrix} = \begin{pmatrix} \beta_{11} & \beta_{12} & \cdots & \beta_{1p} \\ \beta_{21} & \beta_{22} & \cdots & \beta_{2p} \\ \vdots & \vdots & \ddots & \vdots \\ \beta_{m1} & \beta_{m2} & \cdots & \beta_{mp} \end{pmatrix} \in \mathbb{R}^{m \times p}, \quad (79)$$

and $p(\mathbf{B})$ is the penalty function.

The derivative of $f(\mathbf{B})$ with respect to $\mathbf{B}$ is

$$\frac{\partial f}{\partial \mathbf{B}} = \mathbf{M}^T \boldsymbol{\Sigma} \mathbf{M} \mathbf{B} \mathbf{X}^T \mathbf{X} - \mathbf{M}^T \boldsymbol{\Sigma} \tilde{\mathbf{Y}} \mathbf{X} + \frac{\partial p(\mathbf{B})}{\partial \mathbf{B}}. \quad (80)$$

4.3.1 Type I Penalty—When $p(\mathbf{B}) = \alpha \sum_{i=1}^m |\beta_{ip}|$, its derivative with respect to $\mathbf{B}$ is

$$\frac{\partial p(\mathbf{B})}{\partial \mathbf{B}} = \alpha \begin{pmatrix} 0 & \cdots & 0 & \partial |\beta_{1p}| \\ 0 & \cdots & 0 & \partial |\beta_{2p}| \\ \vdots & \ddots & \vdots & \vdots \\ 0 & \cdots & 0 & \partial |\beta_{mp}| \end{pmatrix} = \begin{pmatrix} \mathbf{0}_{m \times (p-1)} & \alpha \frac{\partial \|\boldsymbol{\beta}^p\|_1}{\partial \boldsymbol{\beta}^p} \end{pmatrix}. \quad (81)$$

Denote

$$\begin{aligned} X &= [x^1 \dots x^{p-1} x^p] = [X_1 \ x^p], \\ B &= [\beta^1 \dots \beta^{p-1} \beta^p] = [B_1 \ \beta^p]. \end{aligned}$$

Setting (80) equal to zero gives

$$M^T \Sigma M (B_1 X_1^T + \beta^p x^{pT}) X_1 = M^T \Sigma \tilde{Y} X_1 \quad (82)$$

$$M^T \Sigma M (B_1 X_1^T + \beta^p x^{pT}) x^p + \alpha \frac{\partial \|\beta^p\|_1}{\partial \beta^p} = M^T \Sigma \tilde{Y} x^p. \quad (83)$$

Since $M^T \Sigma M$ is rank deficient², the solution to (82) is not unique. We apply the pseudoinverse of $M^T \Sigma M$ to obtain the minimum-norm solution to (82):

$$B_1 = (M^T \Sigma M)^\dagger (M^T \Sigma \tilde{Y} - M^T \Sigma M \beta^p x^{pT}) X_1 (X_1^T X_1)^{-1}. \quad (84)$$

Substituting (84) into (83) yields

$$\begin{aligned} & M^T \Sigma M \beta^p x^{pT} [I_n - X_1 (X_1^T X_1)^{-1} X_1^T] x^p + \alpha \frac{\partial \|\beta^p\|_1}{\partial \beta^p} \quad (85) \\ &= M^T \Sigma \tilde{Y} [I_n - X_1 (X_1^T X_1)^{-1} X_1^T] x^p. \end{aligned}$$

2. Simple analysis shows that the rank of $M^T \Sigma M$ is $m-1$.

$$\begin{aligned} L_\rho(\{\beta_i\}, \delta_0, \lambda) &= \sum_{i=1}^m \frac{1}{2\sigma_i^2} \sum_{j=1}^n (\tilde{y}_{ij} - x_j^T \beta_i + x_j^T \delta_0)^2 + \alpha \sum_{i=1}^m \|\beta_i\| \quad (73) \\ &+ \lambda^T \left(\frac{1}{\sum_{i=1}^m \frac{1}{\sigma_i^2}} \sum_{i=1}^m \frac{1}{\sigma_i^2} \beta_i - \delta_0 \right) + \frac{\rho}{2} \left\| \frac{1}{\sum_{i=1}^m \frac{1}{\sigma_i^2}} \sum_{i=1}^m \frac{1}{\sigma_i^2} \beta_i - \delta_0 \right\|^2 \end{aligned}$$

$$\begin{aligned}
L_i(\{\beta_i\}, \delta_0, \lambda) &= \frac{1}{2\sigma_i^2} \sum_{j=1}^n (\tilde{y}_{ij} - \mathbf{x}_j^T \beta_i + \mathbf{x}_j^T \delta_0)^2 + \alpha \|\beta_i\| + \lambda^T \frac{1}{\sum_{\ell=1}^m \frac{1}{2} \frac{1}{\sigma_\ell^2}} \beta_i + \frac{\rho}{2} \left\| \frac{1}{\sum_{\ell=1}^m \frac{1}{2} \frac{1}{\sigma_\ell^2}} \beta_i + \frac{1}{\sum_{\ell=1}^m \frac{1}{2} \frac{1}{\sigma_\ell^2}} \sum_{\ell \neq i} \frac{1}{\sigma_\ell^2} \beta_\ell - \delta_0 \right\|^2 \\
&= \frac{1}{2} \beta_i^T \left(\frac{1}{\sigma_i^2} \mathbf{X}^T \mathbf{X} + \frac{\rho}{\sigma_i^4} \frac{1}{\left(\sum_{\ell=1}^m \frac{1}{2} \frac{1}{\sigma_\ell^2} \right)^2} \mathbf{I}_p \right) \beta_i - \beta_i^T \mathbf{v}_i + \alpha \|\beta_i\|
\end{aligned}$$

+ c

(74)

Note that to arrive at (85), we have exploited the fact that $(\mathbf{M}^T \Sigma \mathbf{M})(\mathbf{M}^T \Sigma \mathbf{M})^\dagger \mathbf{M}^T \Sigma = \mathbf{M}^T \Sigma$ which is due to that $\mathbf{M}^T \Sigma \mathbf{M} = \mathbf{M}^T \Sigma$ according to the definition of $\mathbf{M}$ in (37) and definition of the pseudoinverse of a matrix.

Since the coefficient matrix of β^p , i.e., $\mathbf{M}^T \Sigma \mathbf{M} \left(\mathbf{x}^{pT} [\mathbf{I}_n - \mathbf{X}_1 (\mathbf{X}_1^T \mathbf{X}_1)^{-1} \mathbf{X}_1^T] \mathbf{x}^p \right)$ is positive semidefinite, (85) implies that when

$$\begin{aligned}
\alpha &\geq \left\| \mathbf{M}^T \Sigma \tilde{\mathbf{Y}} [\mathbf{I}_n - \mathbf{X}_1 (\mathbf{X}_1^T \mathbf{X}_1)^{-1} \mathbf{X}_1^T] \mathbf{x}^p \right\|_\infty \quad (86) \\
&= \left\| \Sigma \tilde{\mathbf{Y}} [\mathbf{I}_n - \mathbf{X}_1 (\mathbf{X}_1^T \mathbf{X}_1)^{-1} \mathbf{X}_1^T] \mathbf{x}^p \right\|_\infty \\
&= \max_i \left| \frac{1}{\sigma_i^2} \tilde{\mathbf{y}}_i^T [\mathbf{I}_n - \mathbf{X}_1 (\mathbf{X}_1^T \mathbf{X}_1)^{-1} \mathbf{X}_1^T] \mathbf{x}^p \right|,
\end{aligned}$$

where the next to last equality is due to (40) we obtain zero solution.

4.3.2 Type II Penalty—The derivative of $p(\mathbf{B}) = \alpha \sum_{i=1}^m \|\beta_i\|$ with respect to $\mathbf{B}$ is

$$\frac{\partial p(\mathbf{B})}{\partial \mathbf{B}} = \alpha \begin{pmatrix} \frac{\partial \|\boldsymbol{\beta}_1\|}{\partial \boldsymbol{\beta}_1^T} \\ \frac{\partial \|\boldsymbol{\beta}_2\|}{\partial \boldsymbol{\beta}_2^T} \\ \vdots \\ \frac{\partial \|\boldsymbol{\beta}_m\|}{\partial \boldsymbol{\beta}_m^T} \end{pmatrix}, \quad (87)$$

when $\frac{\partial \|\boldsymbol{\beta}_i\|}{\partial \boldsymbol{\beta}_i} = \frac{\boldsymbol{\beta}_i}{\|\boldsymbol{\beta}_i\|}$ if $\boldsymbol{\beta}_i \neq \mathbf{0}$ and $\left\| \frac{\partial \|\boldsymbol{\beta}_i\|}{\partial \boldsymbol{\beta}_i} \right\| \leq 1$ otherwise [27],[28]

Setting (80) equal to zero yields

$$\mathbf{X}^T \mathbf{X} \mathbf{B}^T \mathbf{M}^T \boldsymbol{\Sigma} \mathbf{m}^i - \mathbf{X}^T \tilde{\mathbf{Y}}^T \boldsymbol{\Sigma} \mathbf{m}^i + \alpha \frac{\partial \|\boldsymbol{\beta}_i\|}{\partial \boldsymbol{\beta}_i} = \mathbf{0}_{p \times 1}, \quad (88)$$

for $i=1, 2, \dots, m$, where $\mathbf{m}^i$ is the i -th column of $\mathbf{M}$ in (37) The minimizer to $f(\mathbf{B})$ is a zero matrix when

$$\alpha \geq \max_i \left\| \mathbf{X}^T \tilde{\mathbf{Y}}^T \boldsymbol{\Sigma} \mathbf{m}^i \right\|. \quad (89)$$

Note that

$$\tilde{\mathbf{Y}}^T \boldsymbol{\Sigma} \mathbf{m}^i = \tilde{\mathbf{Y}}^T \boldsymbol{\Sigma} \left(\mathbf{e}_i - \frac{1}{\sum_{\ell=1}^m \frac{1}{\sigma_{\ell}^2}} \frac{1}{\sigma_i^2} \mathbf{I}_m \right) = \tilde{\mathbf{Y}}^T \boldsymbol{\Sigma} \mathbf{e}_i = \frac{1}{\sigma_i^2} \tilde{\mathbf{y}}_i, \quad (90)$$

where the next to last equality is due to (25) Substituting (90) into (89) yields

$$\alpha_{\max} = \max_i \left\| \frac{1}{\sigma_i^2} \mathbf{X}^T \tilde{\mathbf{y}}_i \right\|. \quad (91)$$

4.4 Maximum likelihood estimation of $\{\sigma_i^2\}_{i=1}^m$

To solve for $\{\sigma_i^2\}_{i=1}^m$ consider the negative log-likelihood function with $\{\sigma_i^2\}_{i=1}^m$ being unknown parameters as well:

$$\ell(\beta_0, \{\beta_i\}_{i=1}^m, d, \{\sigma_i^2\}_{i=1}^m) = \sum_{i=1}^m \left[\frac{n}{2} \log(2\pi\sigma_i^2) + \frac{1}{2\sigma_i^2} \sum_{j=1}^n (y_{ij} - \beta_{i0} - \mathbf{x}_j^T \beta_i - d_j)^2 \right]. \quad (92)$$

Taking the partial derivatives of ℓ with respect to d_j and β_{i0} and setting the result to zero, we arrive at

$$d_j = d_j - d_1 = (\bar{y}_{\cdot j}^{(w)} - \bar{y}_{\cdot 1}^{(w)}) - (\mathbf{x}_j - \mathbf{x}_1)^T \bar{\boldsymbol{\beta}}^{(w)}, \quad (93)$$

$$\beta_{i0} = \bar{y}_{i\cdot} - \frac{1}{n} \sum_{j=1}^n \mathbf{x}_j^T \beta_i - \frac{1}{n} \sum_{j=1}^n d_j = \bar{y}_{i\cdot} + \bar{y}_{\cdot 1}^{(w)} - \bar{y}^{(w)} - \mathbf{x}_1^T \bar{\boldsymbol{\beta}}^{(w)}, \quad (94)$$

where to derive the second equality we have exploited assumption (58).

The sum of (93) and (94) gives

$$\beta_{i0} + d_j = \bar{y}_{i\cdot} + \bar{y}_{\cdot j}^{(w)} - \bar{y}^{(w)} - \mathbf{x}_j^T \bar{\boldsymbol{\beta}}^{(w)}. \quad (95)$$

Taking the partial derivative of ℓ with respect to β_i and setting the result to zero, we have

$$\beta_i = \sum_{j=1}^n \mathbf{x}_j y_{ij} - \sum_{j=1}^n \mathbf{x}_j (\beta_{i0} + d_j). \quad (96)$$

Substituting (95) into (96) yields

$$\beta_i = (\mathbf{X}^T \mathbf{X})^{-1} \left[\sum_{j=1}^n \mathbf{x}_j y_{ij} - \frac{1}{\sum_{i=1}^m \frac{1}{\sigma_i^2}} \sum_{i=1}^m \frac{1}{\sigma_i^2} \sum_{j=1}^n \mathbf{x}_j y_{ij} \right] + \bar{\boldsymbol{\beta}}^{(w)}, \quad (97)$$

where $\bar{\boldsymbol{\beta}}^{(w)}$ is defined in (60).

Taking the partial derivative of σ_i^2 and setting the result to zero gives

$$\sigma_i^2 = \frac{1}{n} \sum_{j=1}^n (y_{ij} - \beta_{i0} - \mathbf{x}_j^T \beta_i - d_j)^2. \quad (98)$$

Substituting (95) into (98) yields

$$\sigma_i^2 = \frac{1}{n} \sum_{j=1}^n \left(y_{ij} - \bar{y}_{i.} - \bar{y}_{.j}^{(w)} + \bar{y}^{(w)} - \mathbf{x}_j^T \boldsymbol{\beta}_i + \mathbf{x}_j^T \bar{\boldsymbol{\beta}}^{(w)} \right)^2, \quad (99)$$

where $\bar{y}_{i.}$, $\bar{y}_{.j}^{(w)}$ and $\bar{y}^{(w)}$ are defined in (15), (11) and (18), respectively.

Given initial estimates for $\bar{\boldsymbol{\beta}}^{(w)}$ and $\{\sigma_i^2\}_{i=1}^m$ estimates for $\boldsymbol{\beta}_i$ and σ_i^2 can then be iteratively updated using equations (97), (99), and (60) and until convergence.

After estimating σ_i^2 's, they can then be shrunk (squeezed) toward the common noise variance to obtain robust estimates for σ_i^2 , as done in Section 3.4.

Given initial estimates for $\bar{\boldsymbol{\beta}}^{(w)}$ and $\{\sigma_i^2\}_{i=1}^m$, estimates for $\boldsymbol{\beta}_i$ and σ_i^2 can then be iteratively updated using equations (97), (99), (60) until convergence, as shown in Algorithm 5.

5 EXPERIMENTS

We evaluate the performance of the proposed algorithm (referred to as ELMSeq, short for extended linear model for RNA-seq data analysis). To save space, we only verify the proposed algorithm for the simple regression model (2). We use the 5th percentile to set the tuning parameter α (see Section 3.3).

We compare our method with the state-of-the-art methods for detecting differential gene expression from RNA-seq data: edgeR-robust [20], [29], DESeq2 [18], and limma-voom [16], [17].

5.1 Simulations on Synthetic Data

We simulate RNA-seq data with a total of m 1000 genes and $n = 15$ samples. The data generation is described in Table 1.

We first examine whether the proposed algorithm can accurately estimate the log-fold changes (or slopes) β_i' s. For ease of illustration, we set the true slopes for DE ones as $\beta_i = \pm 2$ instead of $\beta_i \sim \mathcal{N}(\pm 2, 1)$. We start with 300 DE and 700 non-DE genes. Among DE genes 50% are up-regulated while the remaining 50% are down-regulated. The fitted $\{\beta_i\}_{i=1}^m$ using ELMSeq are plotted in Figure 1(a). We see that the estimated slopes are centered around the true ones: the estimated β_i' s of the DE genes are centered around ± 2 , while those of the non-DE genes are close to zero. In Figure 1(b) and Figure 1(c), we increase the percent of up-regulated DE genes to 70% and 90%, respectively. Our method still accurately retrieves all non-zero β_i' s while shrinking all other β_i' s to zero.

In Figure 1(d–f), we increase the number of DE genes to 500, among which 50%, 70% or 90% are up-regulated while others are down-regulated. Our method still achieves accurate estimates. In Figure 1(g–h), we further increase the number of DE genes to 700 among which 50% or 70% are up-regulated, for which our method still achieves accurate estimates when. Only when we simulate with 700 DE genes among which 90% are up-regulated, our method fails to distinguish between DE and non-DE genes since the estimated regression coefficients of the latter are not zero either [Figure 1(i)]. A theoretical explanation of Figure 1(i) has been provided in the supplementary material.

Using a different gene expression measure such as CPM, RPKM or TPM values computed with formulas in (1) yields essentially the same result.

Using the algorithm in Algorithm (1), we estimate the regression coefficient $\hat{\beta}_i$ for each gene i . We decide there is a linear relationship between the predictor variable x_j and the expression data y_{ij} if $\hat{\beta}_i \neq 0$. The larger $|\hat{\beta}_i|$ is, the stronger the relationship. We then sort the genes in descending order of their $|\hat{\beta}_i|$ vary the threshold to construct the receiver operating characteristic (ROC) curve and to calculate the area under the ROC curve (AUC).

The AUCs for DE gene detection using all four methods are summarized in Table 2. We see that the ELMSeg performs better than or comparably to other three methods, regardless of how many genes are differentially expressed and whether they are expressed in a symmetric manner or not. In challenging cases where a large proportion of genes are differentially expressed in an asymmetric manner (e.g., 50% DE genes among which 90% are up-regulated or 70% DE genes among which 70% are up-regulated), the performance gain of the ELMSeg over completing methods is more significant.

In Table 3, we decrease the log-fold change of the DE genes as $\beta_i \sim \mathcal{N}(\pm 0.2, 0.1)$ while keeping all other data generation parameters (including the noise level) the same as those in Table 2. We see that all methods suffer a degradation in AUC performance; but again, the ELMSeg consistently perform better than or comparably to all other methods.

Note that when more samples are available, the performance gain of the ELMSeg over completing methods becomes even more significant. The results for various sample sizes $n=5, 8, 25, 50, 100$ are provided in the supplementary materials (Tables S1–S5 for genes with high expression profiles $\beta_i \sim \mathcal{N}(\pm 2, 1)$ and Tables S6–S10 for genes with low expression profiles $\beta_i \sim \mathcal{N}(\pm 0.2, 0.1)$).

We also performed simulations with the multiple linear regression model in Section 4, and the preliminary results are similar to that obtained for the simple regression model. Note that unlike the simple regression model and type I penalized multiple linear regression model, the type II penalized multiple linear regression model does not allow to define up-and down-regulated genes as multiple regression coefficients are tested simultaneously.

5.2 An application to a real RNA-Seq dataset

We further evaluate our algorithm on a prostate adenocarcinoma (PRAD) RNA-Sequencing dataset published as part of The Cancer Genome Atlas (TCGA) project [30]. The RNA-Seq datasets of 20531 genes from 187 samples were downloaded from the TCGA data portal (<https://tcga-data.nci.nih.gov/tcga>). We desire to identify genes that are associated with pre-operative prostate-specific antigen (PSA), an important risk factor for prostate cancer. The gene expression data were preprocessed by the TCGA consortium. Tissue samples from 333 PRAD patients were sequenced using the Illumina sequencing instruments. The raw sequencing reads were processed and analyzed using the SeqWare Pipeline 0.7.0 and MapSpliceRSEM workflow 0.7 developed by the University of North Carolina, and then aligned to the human reference genome using MapSplice[31]. The gene expression distributions of all samples are normalized to have the same 75th percentile expression values (1,000).

Using the algorithm in Algorithm 1, we obtain the estimated between-sample normalization factors $\hat{d}_j$'s and regression coefficient $\hat{\beta}_i$ for each gene i . We then substitute $\hat{d}_j$'s into model (2), and for each gene i compute the p-value by testing the null hypothesis that the slope of the regression line is equal to zero, i.e., $\beta_i = 0$. We determine a gene is differentially expressed if the p value associated with its linear regression model is less than $0.05/m$. Here the threshold $0.05/m$ is determined using the Bonferroni correction to adjust for multiple significant tests and to achieve a desired family-wise error rate of 0.05. The relations between the sets of differentially expressed genes selected by edgeR, DESeq2, limma-voom and ELMSeg are depicted in Fig. 2.

Nine genes are uniquely detected by ELMSeg: *RIC3*, *ALDH1A2*, *BCL11A*, *CDH3*, *DIRAS3*, *EPHA5*, *CEACAM1*, *PRSS16*, and *AJAPI*. For most of these genes, evidence has also been reported in the literature on their association with prostate cancer. For example, the genes *ALDH1a2* [32] and *CEACAM1* [33] are reported to be tumor suppressors in prostate cancer: underexpression of these genes promote prostate cancer cell proliferation.

Twelve genes are detected by all four methods: *KANK4*, *RHOU*, *TPT1*, *SH2D3A*, *EEF1A1P9*, *ZCWPW1*, *ZNF454*, *RACGAP1*, *PTPLA*, *POC1A*, *AURKA* and *TIMM17A*. The common genes detected by three methods are: six genes *CDK1*, *FAM111B*, *MLF1IP*, *PRC1*, *DTL*, *RAD54B* by edgeR, DESeq2, and limma-voom, three genes *SH3RF2*, *ATCAY* and *PCP4* by edgeR, DESeq2 and ELMSeg, three genes *FERMT1*, *FOXA3* and *LRAT* by edgeR, limma-voom and ELMSeg, and one gene *IPO9* by DESeq2, limma-voom and ELMSeg. For most of these genes, evidence has also been reported in the literature on their association with prostate cancer. For example, the silencing of gene *RHOU* decreases the invasion, proliferation and motility of prostate cancer cells[34].

6 DISCUSSION

A unified statistical model is proposed for joint between sample normalization and DE detection of RNA-seq data. The sample-specific normalization factors are modeled as unknown parameters and jointly estimated together with DE detection. As a result, the

model is robust against normalization errors and is independent of the units (i.e., counts, CPM/RPM, RPKM/FPKM or TPM) in which gene expression levels are summarized.

For the model with a single treatment condition, we introduce the L1 penalty to the linear regression model. The L1 penalty favors sparse solutions (forces some coefficients to be exactly zero). This is desirable since many genes are not differentially expressed. From a Bayesian point of view, the lasso penalty corresponds to a Laplace (double exponential centred at zero) prior over the regression coefficients. By contrast, existing methods do not exploit the sparsity information. We also extend the simple linear regression model to multiple linear regression model to accommodate multiple treatment conditions. Two types of penalty functions are introduced. In the first one only one covariate is of interest while all other covariates are treated as confounding factors. We are interested in testing whether that specific covariate is associated with differential expression. In the second case all covariates are of interest (there are no confounding covariates) and we are interested in testing whether any covariate affects the differential expression of a gene.

Simulation studies show that the proposed methods always perform better than or comparably to existing methods in terms of AUC. The performance gain increases with a larger sample size or higher signal to noise ratio, and is more significant when a large proportion of genes are differentially expressed in an asymmetric manner.

The R codes of the algorithms described in the paper are available for download at <http://www-personal.umich.edu/~jianghui/lr-ADMM/>.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

ACKNOWLEDGMENTS

The authors would like to thank the anonymous reviewers for their valuable comments and suggestions to improve the quality of the paper.

Biographies

Kefei Liu received B.Sc. in mathematics and Ph.D. in Electronic Engineering from Wuhan University in 2006 and City University of Hong Kong in 2013, respectively. He is currently a Postdoctoral Research Fellow at the Center for Computational Biology and Bioinformatics, Indiana University School of Medicine. Before joining IU, he worked as a Postdoctoral Research Associate at The Biodesign Institute of Arizona State University and the Department of Computational Medicine and Bioinformatics of University of Michigan. His

current research interests include machine learning, optimization, tensor decompositions and their applications in biomedical data analysis.

Jieping Ye received the Ph.D. degree in computer science from the University of Minnesota, Twin Cities, MN, USA, in 2005.

He is an Associate Professor of Department of Computational Medicine and Bioinformatics and Department of Electrical Engineering and Computer Science at the University of Michigan, Ann Arbor, MI, USA. His research interests include machine learning, data mining, and biomedical informatics. Dr. Ye has served as Senior Program Committee/Area Chair/Program Committee Vice Chair of many conferences including NIPS, ICML, KDD, IJCAI, ICDM, SDM, ACML, and PAKDD. He serves as a PC Co-Chair of SDM 2015. He serves as an Associate Editor for IEEE TRANSACTIONS ON KNOWLEDGE AND DATA ENGINEERING and IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE, and serves as an Action Editor for Data Mining and Knowledge Discovery. He won the NSF CAREER Award in 2010. His papers have been selected for the outstanding student paper at ICML in 2004, the KDD best research paper honorable mention in 2010, the KDD best research paper nomination in 2011 and 2012, the SDM best research paper runner up in 2013, the KDD best research paper runner up in 2013, and the KDD best student paper award in 2014.

Yang Yang is a Ph.D. candidate at Beihang University. Now he is a visiting student under the supervision of distinguished professor Philip S. Yu at the University of Illinois at Chicago. He got his bachelor's and master's degree from Xidian University. His research interests are social network analysis, machine learning, and complex networks.

Li Shen holds a B.S. degree from Xi'an Jiao Tong University, an M.S. degree from Shanghai Jiao Tong University, and a Ph.D. degree from Dartmouth College, all in Computer Science. He is an Associate Professor of Radiology and Imaging Sciences at Indiana

University School of Medicine. His research interests include medical image computing, bioinformatics, data mining, network science, systems biology, brain imaging genomics, and brain connectomics.

Hui Jiang is an Assistant Professor in the Department of Biostatistics at University of Michigan. He received his Ph.D. in Computational and Mathematical Engineering from Stanford University in 2009. He received his B.S. and M.S. in Computer Science from Peking University. Before joining the University of Michigan in 2011, he was a postdoctoral scholar in the Department of Statistics and Genome Technology Center at Stanford University. He is interested in developing statistical and computational methods for the analysis of large-scale biological data generated using modern high-throughput technologies.

REFERENCES

- [1]. Mortazavi A, Williams BA, McCue K, Schaeffer L, and Wold B, "Mapping and quantifying mammalian transcriptomes by RNA-Seq." *Nat Methods*, vol. 5, no. 7, pp. 621–628, 7 2008. [PubMed: 18516045]
- [2]. Trapnell C, Williams BA, Pertea G, Mortazavi A, Kwan G, van Baren MJ, Salzberg SL, Wold BJ, and Pachter L, "Transcript assembly and quantification by RNA-Seq reveals unannotated transcripts and isoform switching during cell differentiation," *Nat Biotechnol*, vol. 28, no. 5, pp. 511–515, 5 2010. [PubMed: 20436464]
- [3]. Jiang H and Wong WH, "Statistical inferences for isoform expression in RNA-Seq," *Bioinformatics*, vol. 25, no. 8, pp. 1026–1032, 4 2009. [PubMed: 19244387]
- [4]. Salzman J, Jiang H, and Wong WH, "Statistical modeling of RNA-Seq data," *Statistical Science*, vol. 26, no. 1, pp. 62–83, 2011.
- [5]. Wang Z, Gerstein M, and Snyder M, "RNA-Seq: a revolutionary tool for transcriptomics." *Nat Rev Genet*, vol. 10, no. 1, pp. 57–63, 1 2009. [PubMed: 19015660]
- [6]. Dillies M-A, Rau A, Aubert J, Hennequet-Antier C, Jean mougin M, Servant N, Keime C, Marot G, Castel D, Estelle J, Guernec G, Jagla B, Jouneau L, Laloe D, Le Gall C, Schaeffer B, ffer L, Le Crom S, Guedj M, Jaffrezic F, and F. S. C., "A comprehensive evaluation of normalization methods for illumina high-throughput RNA sequencing data analysis," *Brief Bioinform*, vol. 14, no. 6, pp. 671–683, 11 2013. [PubMed: 22988256]
- [7]. Rapaport F, Khanin R, Liang Y, Pirun M, Krek A, Zumbo P, Mason CE, Succi ND, and Betel D, "Comprehensive evaluation of differential gene expression analysis methods for RNA-seq data," *Genome Biology*, vol. 14, no. 9, p. R95, 9 2013. [PubMed: 24020486]
- [8]. Robinson MD and Oshlack A, "A scaling normalization method for differential expression analysis of RNA-seq data," *Genome Biol*, vol. 11, no. 3, p. R25, 2010. [PubMed: 20196867]
- [9]. Bolstad BM, Irizarry RA, Astrand M, and Speed TP, "A comparison of normalization methods for high density oligonucleotide array data based on variance and bias." *Bioinformatics*, vol. 19, no. 2, pp. 185–193, 1 2003. [PubMed: 12538238]
- [10]. Smyth GK, "Limma: linear models for microarray data," in *Bioinformatics and computational biology solutions using R and Bioconductor*. Springer, 2005, pp. 397–420.
- [11]. Bullard JH, Purdom E, Hansen KD, and Dudoit S, "Evaluation of statistical methods for normalization and differential expression in mRNA-Seq experiments." *BMC Bioinformatics*, vol. 11, p. 94, 2010. [PubMed: 20167110]

- [12]. Anders S and Huber W, "Differential expression analysis for sequence count data," *Genome Biol*, vol. 11, no. 10, p. R106, 2010. [PubMed: 20979621]
- [13]. Li B, Ruotti V, Stewart RM, Thomson JA, and Dewey CN, "RNA-Seq gene expression estimation with read mapping uncertainty," *Bioinformatics*, vol. 26, no. 4, pp. 493–500, 2 2010. [PubMed: 20022975]
- [14]. Oshlack A, Wakefield MJ et al., "Transcript length bias in RNA seq data confounds systems biology," *Biol Direct*, vol. 4, no. 1, p. 14, 2009. [PubMed: 19371405]
- [15]. Marioni JC, Mason CE, Mane SM, Stephens M, and Gilad Y, "RNA-seq: an assessment of technical reproducibility and comparison with gene expression arrays," *Genome research*, vol. 18, no. 9, pp. 1509–1517, 2008. [PubMed: 18550803]
- [16]. Law CW, Chen Y, Shi W, and Smyth GK, "Voom: precision weights unlock linear model analysis tools for RNA-seq read counts," *Genome Biol*, vol. 15, no. 2, p. R29, 2014. [PubMed: 24485249]
- [17]. Ritchie ME, Phipson B, Wu D, Hu Y, Law CW, Shi W, and Smyth GK, "limma powers differential expression analyses for RNA-sequencing and microarray studies," *Nucleic Acids Research*, vol. 43, no. 7, p. e47, 1 2015. [PubMed: 25605792]
- [18]. Love MI, Huber W, and Anders S, "Moderated estimation of fold change and dispersion for RNA-seq data with DESeq2," *Genome Biology*, vol. 15, no. 12, p. 550, 2014. [PubMed: 25516281]
- [19]. Jiang H and Salzmann J, "A penalized likelihood approach for robust estimation of isoform expression," *Statistics and Its Interface*, vol. 8, no. 4, pp. 437–445, 2015. [PubMed: 27239250]
- [20]. Zhou X, Lindsay H, and Robinson MD, "Robustly detecting differential expression in RNA sequencing data using observation weights," *Nucleic acids research*, vol. 42, no. 11, pp. e91–e91, 2014. [PubMed: 24753412]
- [21]. Boyd S, Parikh N, Chu E, Peleato B, and Eckstein J, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Foundations and Trends® in Machine Learning*, vol. 3, no. 1, pp. 1–122, 2011.
- [22]. Jiang H and Zhan T, "Unit-free and robust detection of differential expression from RNA-Seq data," *Statistics in Biosciences*, vol. 9, no. 1, pp. 178–199, 2017.
- [23]. Tibshirani R, "Regression shrinkage and selection via the lasso," *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 267–288, 1996.
- [24]. Ji H and Wong WH, "TileMap: create chromosomal map of tiling array hybridizations," *Bioinformatics*, vol. 21, no. 18, pp. 3629–3636, 2005. [PubMed: 16046496]
- [25]. Ji H and Liu XS, "Analyzing 'omics data using hierarchical models," *Nature biotechnology*, vol. 28, no. 4, pp. 337–340, 2010.
- [26]. Smyth G, "Statistical applications in genetics and molecular biology," *Linear models and empirical Bayes methods for assessing differential expression in microarray experiments*, 2004.
- [27]. Friedman J, Hastie T, and Tibshirani R, "Regularization paths for generalized linear models via coordinate descent," *Journal of statistical software*, vol. 33, no. 1, p. 1, 2010. [PubMed: 20808728]
- [28]. Yuan M and Lin Y, "Model selection and estimation in regression with grouped variables," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 68, no. 1, pp. 49–67, 2006.
- [29]. Robinson MD, McCarthy DJ, and Smyth GK, "edgeR: a bioconductor package for differential expression analysis of digital gene expression data," *Bioinformatics*, vol. 26, pp. 139–140, 1 2010. [PubMed: 19910308]
- [30]. Network CGAR, "The molecular taxonomy of primary prostate cancer," *Cell*, vol. 163, pp. 1011–1025, 11 2015. [PubMed: 26544944]
- [31]. Wang K, Singh D, Zeng Z, Coleman SJ, Huang Y, Savich GL, He X, Mieczkowski P, Grimm SA, Perou CM, MacLeod JN, Chiang DY, Prins JF, and Liu J, "Mapsplice: accurate mapping of RNA-seq reads for splice junction discovery," *Nucleic acids research*, vol. 38, p. e178, 10 2010. [PubMed: 20802226]
- [32]. Kim H, Lapointe J, Kaygusuz G, Ong DE, Li C, van de Rijn M, Brooks JD, and Pollack JR, "The retinoic acid synthesis gene ALDH1a2 is a candidate tumor suppressor in prostate cancer," *Cancer research*, vol. 65, no. 18, pp. 8118–8124, 2005. [PubMed: 16166285]

- [33]. Busch C, Hanssen TA, Wagener C, and Öbrink B, “Down regulation of CEACAM1 in human prostate cancer: correlation with loss of cell polarity, increased proliferation rate, and gleason grade 3 to 4 transition,” *Human pathology*, vol. 33, no. 3, pp. 290–298, 2002. [PubMed: 11979369]
- [34]. Alinezhad S, Väinänen R-M, Mattsson J, Li Y, Tallgrén T, Ochoa NT, Bjartell A, Åkerfelt M, Taimen P, Boström PJ et al., “Validation of novel biomarkers for prostate cancer progression by the combination of bioinformatics, clinical and functional studies,” *PloS one*, vol. 11, no. 5, p. e0155901, 2016. [PubMed: 27196083]

Figure 1:

Estimated β_i in the simple linear regression model from simulated RNASeq data, where the number of genes is $m=1000$ and number of samples is $n=15$. The number of DE genes varies from 300 to 700, and the percentage of up-regulated DE genes varies from 50% to 90%. Along the horizontal axis, from left to right: up-regulated genes ($\beta_i=2$), down-regulated genes ($\beta_i=-2$) and non-DE genes ($\beta_i=0$).

Figure 2: Venn diagram showing the relation between the set of differentially expressed genes detected by edgeR, DESeq2, limma-voom and ELMSeg.

Table 1:

Synthetic data generation process and parameters

$\ell_i \sim 2\text{unif}(5, 10)$	gene length of gene i
$\beta_{i0} \sim \mathcal{N}(0, 1)$	other log scaling factors of gene i
$\beta_i = 0$	log-fold change for non-DE genes
$\beta_i \sim \mathcal{N}(2, 1)$	log-fold change for up-regulated DE genes
$\beta_i \sim \mathcal{N}(-2, 1)$	log-fold change for down-regulated DE genes
$x_j \sim \mathcal{N}(0, 1)$	condition data of sample j
$N_j \sim \text{unif}(2, 3) \times 10^6$	library size of sample j
$d_j \sim \mathcal{N}(0, 1)$	other log scaling factors of sample j
$\mu_{i,j} = N_j \frac{\ell_i}{\sum_{i=1}^m \ell_i} e^{\beta_{i0} + \beta_i x_i + d_j}$	expected RNA-seq read counts of gene i from sample j
$c_{i,j} = \left\lfloor \frac{\log \mathcal{N}(\mu_{i,j}, 0.1)}{e} \right\rfloor$	read counts
$y_{ij} = \log c_{ij}$	log-transformed gene expression

Table 2:

AUC comparison of edgeR-robust, DESeq2, limma voom and ELMSeq in log-normally distributed data. Number of samples: $n=15$, log-fold change for DE genes: $\beta_i \sim \mathcal{N}(\pm 2, 1)$, and noise level: $\sigma_i = 0.1$. The table shows the percent of DE genes (DE %), percent of up-regulated genes among the DE genes (Up %), as well as the mean AUCs for all four methods measured using 10 simulated replicates. The standard errors of the mean AUCs are given in parentheses.

DE (%)	Up (%)	edgeR	DESeq2	voom	ELMSeq
10	50	0.9903 (0.0016)	0.6068 (0.0807)	0.991 (0.0018)	0.9914 (0.0017)
10	70	0.9935 (0.0021)	0.4527 (0.0638)	0.9941 (0.0021)	0.9943 (0.0021)
10	90	0.9869 (0.0028)	0.6878 (0.0637)	0.9875 (0.0024)	0.9897 (0.0022)
30	50	0.9898 (0.001)	0.5508 (0.0883)	0.99 (0.001)	0.99 (0.001)
30	70	0.9891 (0.0014)	0.7946 (0.064)	0.9897 (0.0014)	0.991 (0.0011)
30	90	0.9788 (0.0023)	0.6114 (0.0805)	0.9796 (0.0022)	0.9795 (0.0014)
50	50	0.9917 (8e-04)	0.429 (0.0797)	0.9916 (8e-04)	0.9917 (8e-04)
50	70	0.9748 (0.0026)	0.4923 (0.081)	0.9754 (0.0026)	0.9826 (0.0015)
50	90	0.8717 (0.0133)	0.4697 (0.0667)	0.8801 (0.0119)	0.9662 (0.002)
70	50	0.9907 (9e-04)	0.5572 (0.1027)	0.9915 (8e-04)	0.9923 (7e-04)
70	70	0.8564 (0.018)	0.5307 (0.0588)	0.8696 (0.0148)	0.9591 (0.0034)
70	90	0.3375 (0.0108)	0.4808 (0.0192)	0.3204 (0.0154)	0.4718 (0.0124)

Table 3:

AUC comparison of edgeR-robust, DESeq2, limma voom and ELMSeq in log-normally distributed data. The data generation parameters are the same as those in Table 2 except that the log-fold changes for DE genes decrease to: $\beta_i \sim \mathcal{N}(\pm 0.2, 0.1)$.

DE (%)	Up (%)	edge	DESeq2	voom	ELMSeq
10	50	0.8055 (0.0089)	0.5241 (0.0142)	0.8224 (0.0095)	0.8232 (0.0095)
10	70	0.8086 (0.009)	0.4846 (0.0126)	0.8212 (0.0095)	0.8234 (0.0101)
10	90	0.7867 (0.0084)	0.5078 (0.0084)	0.7955 (0.0104)	0.8024 (0.0106)
30	50	0.8087 (0.005)	0.497 (0.0119)	0.8158 (0.0054)	0.8157 (0.0054)
30	70	0.7848 (0.0052)	0.5471 (0.0211)	0.7949 (0.0052)	0.8013 (0.0055)
30	90	0.7398 (0.0059)	0.5329 (0.0181)	0.7505 (0.0059)	0.773 (0.0054)
50	50	0.8143 (0.0061)	0.4931 (0.0137)	0.8265 (0.0049)	0.8268 (0.0051)
50	70	0.7611 (0.0054)	0.5061 (0.0155)	0.7704 (0.0054)	0.7752 (0.0056)
50	90	0.6451 (0.006)	0.5017 (0.0102)	0.6503 (0.0059)	0.6793 (0.0025)
70	50	0.8149 (0.0022)	0.5231 (0.0273)	0.8261 (0.003)	0.8267 (0.0028)
70	70	0.7271 (0.0074)	0.5093 (0.01)	0.7354 (0.0086)	0.7388 (0.0083)
70	90	0.5449 (0.0066)	0.5158 (0.0089)	0.5505 (0.0081)	0.5434 (0.0069)