

HMMSEQ: A HIDDEN MARKOV MODEL FOR DETECTING DIFFERENTIALLY EXPRESSED GENES FROM RNA-SEQ DATA

BY SHIQI CUI*, SUBHARUP GUHA*,¹, MARCO A. R. FERREIRA^{†,2}
AND ALLISON N. TEGGE[†]

*University of Missouri** and *Virginia Tech*[†]

We introduce *hmmSeq*, a model-based hierarchical Bayesian technique for detecting differentially expressed genes from RNA-seq data. Our novel *hmmSeq* methodology uses hidden Markov models to account for potential co-expression of neighboring genes. In addition, *hmmSeq* employs an integrated approach to studies with technical or biological replicates, automatically adjusting for any extra-Poisson variability. Moreover, for cases when paired data are available, *hmmSeq* includes a paired structure between treatments that incorporates subject-specific effects. To perform parameter estimation for the *hmmSeq* model, we develop an efficient Markov chain Monte Carlo algorithm. Further, we develop a procedure for detection of differentially expressed genes that automatically controls false discovery rate. A simulation study shows that the *hmmSeq* methodology performs better than competitors in terms of receiver operating characteristic curves. Finally, the analyses of three publicly available RNA-seq data sets demonstrate the power and flexibility of the *hmmSeq* methodology. An R package implementing the *hmmSeq* framework will be submitted to CRAN upon publication of the manuscript.

1. Introduction. RNA-seq has revolutionized the study of gene expression. RNA-seq success may be attributed to its low noise, high-throughput and ability to interrogate allele-specific expression and isoforms [Auer, Srivastava and Doerge (2012), Zhao et al. (2014)]. Most RNA-seq studies aim to identify differentially expressed (DE) genes between samples corresponding to different treatments or biological conditions, for example, cancer tissue versus normal tissue, genetically engineered animals versus control animals, or patients exposed to two or more kinds of treatments. These differentially expressed genes usually form the starting point of more extensive studies such as integration of expression data with transcription factor binding [Karlebach and Shamir (2008)], RNA interference [Pe'er and Hacohen (2011)] and DNA methylation [Louhimo and Hautaniemi (2011)], all of which can lead to a better understanding of regulatory mechanisms. Currently available methods for RNA-seq data analysis assume that differential expression of genes occurs independent of the genomic loci of each gene [Auer and

Received February 2014; revised January 2015.

¹Supported by NSF Award DMS-09-06734 and by NIH Award P01CA134294.

²Supported by NSF Award DMS-09-07064.

Key words and phrases. Bayesian hierarchical model, first order dependence, next-generation sequencing, overdispersion, serial correlation.

Doerge (2011), Robinson and Smyth (2007, 2008), Hardcastle and Kelly (2010), Robinson, McCarthy and Smyth (2010), Si and Liu (2013)]. However, the literature contains evidence that neighboring genes on the chromosome tend to be co-expressed [Caron et al. (2001), Michalak (2008), Singer et al. (2005)]. To account for and take advantage of this potential co-expression, here we introduce *hmmSeq*.

Our *hmmSeq* framework incorporates potential co-expression by modeling differential expression across the genome using hidden Markov models (HMM). In the *hmmSeq* framework that we propose, neighboring gene co-expression may occur in two ways: in differential expression across treatments and in mean expression magnitude. Thus, we model gene differential expression across treatments using an HMM with three states: not differentially expressed, under- or over-expressed. This HMM takes advantage of the potential co-differential expression by borrowing information across neighboring genes on the chromosome. In addition, we model gene mean expression magnitude with an HMM with two states: low expression and high expression. The latter HMM borrows information across the genome to increase estimation precision of the mean expression magnitude of each gene. As we show in the simulation study in Section 4, the use of information both from neighboring genes and across the genome increases detection power and reduces false discovery.

The existing methods for RNA-seq data analysis do not account for the potential co-expression of neighboring genes. Robinson and Smyth (2007, 2008) use the negative binomial distribution to model over-dispersed data through dispersion parameters. Specifically, Robinson and Smyth (2008) assume a common dispersion parameter across all tags (or genes), whereas Robinson and Smyth (2007) assume tag-wise (or gene-wise) dispersion parameters. To estimate these dispersion parameters, they assume a Gaussian hierarchical hyperprior that is estimated using empirical Bayes. After that, the gene-wise dispersion parameters are estimated by maximum weighted likelihood. This method is implemented in *edgeR* [Robinson, McCarthy and Smyth (2010)]. The *baySeq* method of Hardcastle and Kelly (2010) is an empirical Bayes approach that is also based on the negative binomial distribution. An empirically determined prior distribution is derived from the entire data set, and rather than producing significance values, this method calculates posterior probabilities of multiple models of differential expression, ranking the genes by the model probabilities. Blekhnman et al. (2010) analyze RNA-seq data by a Poisson generalized linear mixed-effect model, which explains inter-individual variability through the inclusion of a random individual-specific effect. Data are fitted under the null and alternative models gene by gene, then a likelihood ratio test is conducted to compute p -values, and the false discovery rate (FDR; defined as the proportion of incorrect calls among the genes declared as DE) is controlled by the method of Storey and Tibshirani (2003). Auer and Doerge (2011) have proposed the *two-stage Poisson model* (TSPM) which assumes data contain both overdispersed and nonoverdispersed genes. This technique seeks to reduce FDR by first

separating the overdispersed genes from the nonoverdispersed genes, and then fitting separate models to compute the p -values. Benjamini and Hochberg (1995) FDR controlling is applied on each set of p -values to identify DE genes. Si and Liu (2013) developed a test for the hypothesis that the log fold change belongs to a subset of the real line. By assuming parameters under a null and alternative hypothesis come from different distributions, they estimate this mixture distribution from the data, then the test statistic is obtained as the ratio of unconditional probability from null parameter space over unconditional probability from full space. All these previous methods assume that the genes are conditionally independent. However, the exploratory data analysis we present in Section 2.2 suggests dependence among neighboring genes. Our hmmSeq method addresses this dependence.

Our hmmSeq framework may also accommodate the case when there is no dependence among the expression of neighboring genes. Specifically, HMMs include as particular cases mixture models. In particular, the number of components of the mixture model will be the same as the number of states in the HMM. Thus, when there is no co-differential expression, the result will be a mixture model with three components that correspond to a gene being not differentially expressed, under- or over-expressed. Likewise, when there is no dependence in mean expression magnitude among neighboring genes, the resulting mixture model will have two components, one component for low expression genes and another for high expression genes. Note that the proportion of genes in each component and the parameters of the generating model for each component will be estimated from the data. Thus, even without neighboring genes dependence, the hmmSeq framework will still borrow information across the genome to learn about each of the mixture components and, as a result, increase estimation precision and detection power.

We model extra-Poisson variability in an indirect manner. If the experiment contains technical replicates (i.e., samples from the same subject), then the literature provides evidence that the RNA-seq counts are Poisson distributed [Bullard et al. (2010), Marioni et al. (2008)]. On the other hand, if the experiment contains biological replicates, then the RNA-seq counts will have extra-Poisson variability [Langmead et al. (2010), Robinson and Smyth (2007)]. This extra-Poisson variability may be a result of across-subjects variability or slight differences in the experimental conditions when the samples were taken or analyzed. While the RNA-seq literature usually uses the negative binomial distribution to model the extra-Poisson variability, another way to deal with this extra-variability is through the use of the Poisson distribution together with random effects. We prefer the latter because it provides a framework that can flexibly deal with known sources of extra-Poisson variability such as, for example, biological variation among subjects. In the case of paired data considered in Section 5.3, we deal with the biological variation by including for each gene subject-specific random effects. Moreover, for nonpaired data we implicitly deal with the subject-specific random effects (and any other source of extra-Poisson variability) by assuming that for nondifferentially expressed genes the differential treatment effect parameter may come from a

normal distribution centered at zero. Hence, for nondifferentially expressed genes the differential treatment effect parameter is a sum of the random effects of subjects and other hidden sources. In addition to facilitating the implementation of our HMM framework, this aspect of our model increases robustness with respect to hidden unforeseen sources of variability.

We have investigated in three fronts the practical usefulness and adequacy of HMMs and mixture components models for RNA-seq data analysis. First, we have performed an exploratory data analysis presented in Section 2.2 that studies for two real RNA-seq data sets the empirical statistical properties of preliminary estimates of differential expression parameters and mean expression magnitude parameters. Two patterns emerge from this exploratory data analysis: the existence of three differential expression states and of two mean expression magnitude states; and a possible dependence across neighboring genes. Second, we have performed a simulation study that considers all four possible combinations of HMMs and mixture components models for differential expression and mean expression magnitude. This simulation study compares the performance of our *hmmSeq* framework with competing RNA-seq analysis methodologies. In all four possible cases, our *hmmSeq* framework beats the competing methods in terms of receiver operating characteristic curves. Finally, we have used the deviance information criterion (DIC) [Spiegelhalter et al. (2002)] to decide among the four possible combinations of HMMs and mixture components models what is the most adequate model for each of three real RNA-seq data sets. Our use of the DIC is justified by its good performance in a simulation study presented in Section 4. The DIC indicates dependence across neighboring genes for two of the three data sets. Therefore, in this paper we provide further evidence that for some biological processes neighboring genes on the chromosome tend to be co-expressed.

We take a full Bayesian analysis approach and develop a Markov chain Monte Carlo algorithm to exploit the posterior distribution of the model parameters. To simulate the differential effects and the mean effect magnitudes, we develop an efficient Metropolis–Hastings algorithm for hidden Markov models. In addition, we use the output of the MCMC algorithm to identify differentially expressed genes while controlling for false discovery rate. Specifically, we use a Bayesian approach for controlling the FDR level proposed by Newton et al. (2004) and further studied by Müller, Parmigiani and Rice (2007). We demonstrate the advantages and benefits of our *hmmSeq* methodology by analyzing three RNA-seq data sets. The first data set [Marioni et al. (2008)] consists of five technical replicates each of a kidney and liver RNA sample. The second data set [Zeng et al. (2012)] consists of six biological replicates extracted from two regions, frontal pole and hippocampus, of normal human brains. Finally, the third data set [Henn et al. (2013)] consists of paired B-cell samples data of day 0 (before vaccination) and day 7 (post-vaccination) for 3 pre-vaccinated subjects. Therefore, we demonstrate the power and flexibility of the *hmmSeq* methodology on three types of RNA-seq data: technical replicates, biological replicates and paired samples.

The remainder of the paper is organized as follows. Section 2 describes the details of the `hmmSeq` model and informally demonstrates the necessity for hidden Markov models in Section 2.2. Section 3 describes the posterior inference procedure based on Markov chain Monte Carlo (MCMC) techniques and the procedure for identification of DE genes. Section 4 uses simulated data to demonstrate the effectiveness of `hmmSeq` relative to well-known techniques for RNA-seq. The Marioni et al. (2008), Zeng et al. (2012) and Henn et al. (2013) data sets are analyzed in Sections 5.1, 5.2 and 5.3. In all cases, the results are compared and contrasted with those of existing approaches to demonstrate the success of `hmmSeq`. A functional analysis of the detected sets of DE genes provides further evidence of the reliability of the procedure. An R package implementing the `hmmSeq` framework will be submitted to CRAN upon publication of the manuscript.

2. A Bayesian hierarchical model for RNA-seq data. We focus on two-treatment comparisons. For a given chromosome c , let Y_{ijkc} denote the integer-valued gene read of the k th replicate of gene i under treatment j , for gene $i = 1, 2, \dots, I_c$, treatment $j = 1, 2$, and replicate $k = 1, 2, \dots, K_j$ on chromosome c . The genes are sequentially arranged so that consecutive indices correspond to neighboring genes on the chromosome. We assume that

$$(2.1) \quad \begin{aligned} Y_{ijkc} &\stackrel{\text{indep}}{\sim} \text{Poisson}(\lambda_{ijkc}) && \text{where} \\ \log(\lambda_{i1kc}) &= \beta_{ic} - \Delta_{ic} + \rho_{1k} && \text{and} \\ \log(\lambda_{i2kc}) &= \beta_{ic} + \Delta_{ic} + \rho_{2k}, \end{aligned}$$

where β_{ic} denotes the mean expression magnitude of gene i and $2\Delta_{ic}$ denotes the log-fold change between the treatments. The treatment-specific replicate effects are represented by ρ_{jk} . We observe that the treatments are a priori interchangeable in equation (2.1). The *differential treatment effect* Δ_{ic} for gene i is key because it determines the relative expression levels of the treatments for the gene. That is, Δ_{ic} determines whether treatment 2 is over-, under-, or nondifferentially expressed with respect to treatment 1.

To model the mean expression magnitude, the possible dependence among the β_{ic} 's of neighboring genes on a chromosome is modeled using either a two-component finite mixture model [Frühwirth-Schnatter (2006), Titterton, Smith and Makov (1985)] or a stationary two-state hidden Markov model [MacDonald and Zucchini (1997), Rabiner (1989)]. The latent *average expression state* s_{ic} determines whether the expression of gene i , averaged over treatments and replicates, is "small" ($s_{ic} = 1$) or "large" ($s_{ic} = 2$). In the absence of differential treatment and replicate effects, the two levels of this categorical variable correspond, respectively, to low and high reads for the genes. Conditional on s_{ic} , the average expression β_{ic} is normally distributed:

$$\beta_{ic|s_{ic}} \stackrel{\text{indep}}{\sim} N(\mu_{s_{ic}c}, \sigma_{s_{ic}c}^2)$$

with $\mu_{1c} < \mu_{2c}$. The latent states s_{1c}, \dots, s_{I_c} follow either a finite mixture model (FMM) with probability vector $\mathbf{P}_c = (p_{1c}, p_{2c})$ or a hidden Markov model (HMM) with stationary transition probability matrix $\mathbf{A}_c = ((a_{utc}))_{2 \times 2}$ with the row sums $\sum_{t=1,2} a_{utc} = 1$ for $u = 1, 2$. We denote the two-component FMM by \mathcal{F}_{2c} and the two-state HMM by \mathcal{H}_{2c} , assuming independent, noninformative priors for its dispersion parameters: $p(\sigma_{uc}^2) \propto \sigma_{uc}^{-2}$ for $u = 1, 2$.

To model differential expression, the differential effects $\Delta_{1c}, \dots, \Delta_{I_c}$ of the genes are modeled either by a finite mixture model (FMM) \mathcal{F}_{3c} with probability vector $\mathbf{Q}_c = (q_{1c}, q_{2c}, q_{3c})$ or by a three-state stationary HMM denoted by \mathcal{H}_{3c} ; the matrix of transition probabilities is denoted by $\mathbf{B}_c = ((b_{vtc}))_{3 \times 3}$ with row sums $\sum_{t=1}^3 b_{vtc} = 1$ for $v = 1, 2, 3$. With latent *differential states* h_{1c}, \dots, h_{I_c} taking values in $\{1, 2, 3\}$, the values correspond, respectively, to the gene-specific under-, nondifferential-, and over-expression of treatment 2 relative to treatment 1. Given the state h_{ic} , differential effect Δ_{ic} is distributed as

$$(2.2) \quad \Delta_{ic} | h_{ic} \sim \begin{cases} N(\phi_{1c}, \tau_{1c}^2), & \text{if } h_{ic} = 1 \text{ (under-expressed),} \\ N(0, \tau_{2c}^2), & \text{if } h_{ic} = 2 \text{ (nondifferentially-expressed),} \\ N(\phi_{3c}, \tau_{3c}^2), & \text{if } h_{ic} = 3 \text{ (over-expressed),} \end{cases}$$

where $\phi_{1c} < 0$ and $\phi_{3c} > 0$. Thus, for each chromosome, h_{1c}, \dots, h_{I_c} are the parameters of interest because they identify the set of DE genes.

We observe that the latent states of both FMM \mathcal{F}_{2c} and HMM \mathcal{H}_{2c} are nonexchangeable, being associated with particular biological conditions. The priors for the state parameters are designed to reflect this and also to prevent label switching [Scott (2002)]. Specifically, the mean parameters, μ_{1c} and μ_{2c} are assigned the prior $p(\mu_{1c}, \mu_{2c}) \propto 1_{\{\mu_{1c} \leq \mu_{2c} - \delta\}}$ where $\delta > 0$ is a predetermined constant. The fact that δ is strictly positive guarantees that $\mu_{1c} < \mu_{2c}$ and the two states are identifiable.

For the same reason, for the FMM \mathcal{F}_{3c} and HMM \mathcal{H}_{3c} , we assume that $p(\phi_{1c}, \phi_{3c}) \propto 1_{\{\phi_{1c} < u_1, \phi_{3c} > l_3\}}$, where $u_1 < 0$ and $l_3 > 0$ are prespecified constants that can be chosen as follows. The log-fold change between the over- and under-expressed categories is at least $(l_3 - u_1)$. From a practical standpoint, in order to distinguish between these two categories, it is reasonable to assume that the ratios of their associated Δ_{ic} 's exceed 2. Because of this, we symmetrically set $u_1 = -(\log 2)/2$ and $l_3 = (\log 2)/2$. To further facilitate inferences of the state-specific parameters, informative conjugate priors are assigned to τ_{1c}^2, τ_{2c}^2 and τ_{3c}^2 .

2.1. *Paired data analysis.* Our hmmSeq framework may also accommodate paired data, that is, the case when each subject undergoes each of the treatments. Here we describe the minor changes needed for that purpose. For a given chromosome c , let Y_{ijk} denote the gene read of the k th subject of gene i under treatment j , for subject $k = 1, \dots, K$. Obviously, because of the paired data structure

there exists dependence between observations on the same subject. To account for this dependence, we assume that

$$(2.3) \quad \begin{aligned} Y_{ijkc} &\overset{\text{indep}}{\sim} \text{Poisson}(\lambda_{ijkc}) \quad \text{where} \\ \log(\lambda_{i1kc}) &= \beta_{ic} - \Delta_{ic} + \varepsilon_{ikc} + \rho_{1k} \quad \text{and} \\ \log(\lambda_{i2kc}) &= \beta_{ic} + \Delta_{ic} + \varepsilon_{ikc} + \rho_{2k}, \end{aligned}$$

with $\varepsilon_{ikc} \sim N(0, \sigma_\varepsilon^2)$ denoting the subject-specific random effects. The other parameters in the paired-data model have the same interpretations as in equation (2.1).

2.2. Exploratory data analysis. We have performed an exploratory data analysis (EDA) to verify some of the `hmmSeq` model assumptions for the [Marioni et al. \(2008\)](#) and [Zeng et al. \(2012\)](#) data sets. Both data sets possess the feature described by [Bullard et al. \(2010\)](#); for libraries under each treatment, 5% of the genes account for over 50% of the total library size, and 10% of the genes account for over 60% of the total library size. Focusing only on those genes whose reads exceeded nine for all libraries, and ignoring the replicate effects, we preprocessed the raw counts using the upper-quartile normalizing technique of [Bullard et al. \(2010\)](#).

For the genes $i = 1, 2, \dots, I_c$ of each chromosome, we have computed preliminary estimates of the expression magnitude β_{ic} and differential effect Δ_{ic} by treating these parameters as the fixed effects in a Poisson regression model. Figure 1 displays graphical summaries of these estimates for a few chromosomes of the [Marioni et al. \(2008\)](#) data set. The results were similar for the other chromosomes. The density plot for the β_i 's in Figure 1(a) and for the Δ_i 's in Figure 1(b) are indicative of mixture of densities representations for these parameters. Further, data analysis in Section 5.1 will select the dependence structures of β_i 's and Δ_i 's by DIC model selection.

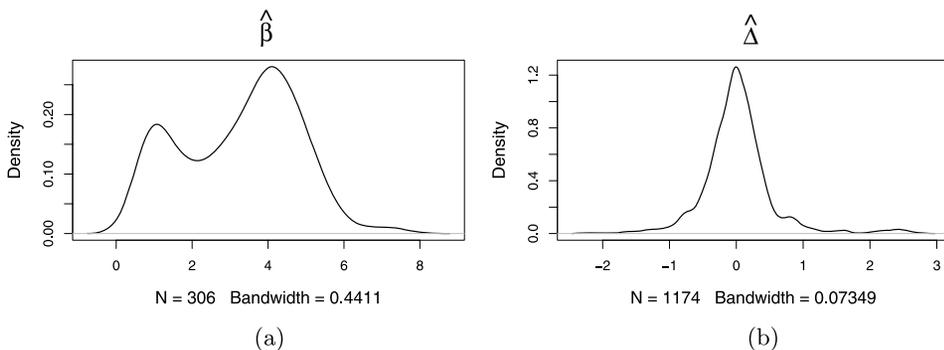


FIG. 1. [Marioni et al. \(2008\)](#) data set—density plots of preliminary estimates for the expression magnitudes β and the differential expression effects Δ . (a) Density plot of $\hat{\beta}_i$ for chromosome 13, (b) density plot of $\hat{\Delta}_i$ for chromosome 19.

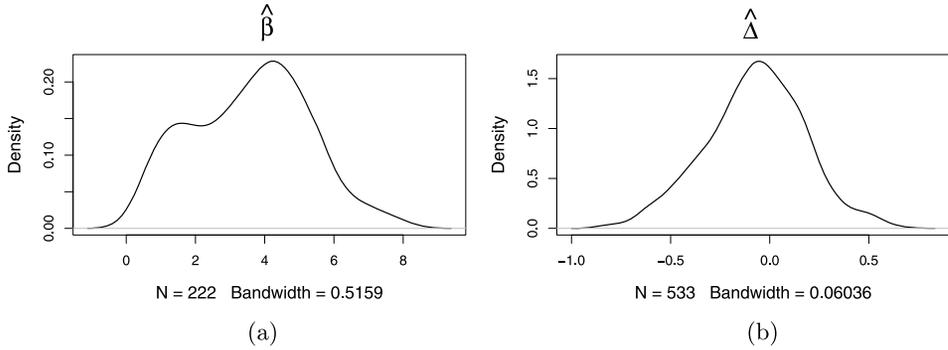


FIG. 2. Zeng et al. (2012) data set—density plots of preliminary estimates for the expression magnitudes β and the differential expression effects Δ . (a) Density plot of $\hat{\beta}_i$ for chromosome 13, (b) density plot of $\hat{\Delta}_i$ for chromosome 9.

For the Zeng et al. (2012) data set, Figure 2 displays density plots for preliminary estimates of the model parameters and also reveals a similar pattern as the Marioni et al. (2008) data set. The above analysis, together with data analysis in Section 5.2, suggests the need for mixture models, with first order dependence to model the expression magnitudes β_{ic} and differential effects Δ_{ic} , justifying the use of the hidden Markov models \mathcal{H}_{2c} and \mathcal{H}_{3c} in the hmmSeq method.

3. Posterior inference. We investigate the posterior distribution of the chromosome-specific parameters using Markov chain Monte Carlo (MCMC) methods. Gibbs sampling cannot be applied to generate the parameters in equation (2.1) because the Poisson likelihood function is not conjugate to the normal priors of the parameters. Consequently, we apply the Laplace approximation [e.g., Chib and Greenberg (1994), Zeger and Karim (1991)] to generate proposed updates for the equation (2.1) parameters. The proposals are accepted or rejected by a Metropolis–Hastings probability to compensate for the use of an approximation instead of the Poisson distribution. This guarantees the convergence of the Markov chain to the posterior distribution of the hmmSeq model. We analyze each chromosome separately and, for simplicity of notation, in this section we omit the chromosome index c .

3.1. Metropolis–Hastings algorithm for HMMs. In this section we present a Metropolis–Hastings algorithm for the simulation of a general latent process $\{\theta_i, i = 1, \dots, I\}$ that follows an HMM \mathcal{H}_m . We use this algorithm in Section 3.2 to simulate the expression magnitude β and the differential effects Δ . Under a Laplace approximation, the *working values* of the read counts are defined as

$$(3.1) \quad w_{ijk} = \log(\lambda_{ijk}) + \frac{Y_{ijk} - \lambda_{ijk}}{\lambda_{ijk}}.$$

These working values have an approximate normal distribution, specifically $w_{ijk} \stackrel{\text{approx}}{\sim} N(\log(\lambda_{ijk}), 1/\lambda_{ijk})$.

For a more general case, suppose θ_i is the parameter of interest, and its value at the previous MCMC iteration was $\theta_i^{(\text{old})}$. Then, the Laplace approximation (3.1) gives us $w_{ijk} = \log(\lambda_{ijk}^*) + (Y_{ijk} - \lambda_{ijk}^*)/\lambda_{ijk}^* \stackrel{\text{approx}}{\sim} N(\log(\lambda_{ijk}), 1/\lambda_{ijk}^*)$, where $\log(\lambda_{ijk}) = \xi_{ijk} + z_{ijk} \cdot \theta_i$ and $\log(\lambda_{ijk}^*) = \xi_{ijk} + z_{ijk} \cdot \theta_i^{(\text{old})}$.

Let the vectors $\mathbf{w}_i = (w_{i11}, w_{i12}, \dots, w_{i2K_2})'$, $\boldsymbol{\lambda}_i = (\lambda_{i11}, \lambda_{i12}, \dots, \lambda_{i2K_2})'$ and $\boldsymbol{\lambda}_i^* = (\lambda_{i11}^*, \lambda_{i12}^*, \dots, \lambda_{i2K_2}^*)'$. Then $\mathbf{w}_i \stackrel{\text{approx}}{\sim} N(\log(\boldsymbol{\lambda}_i), \text{Diag}(1/\boldsymbol{\lambda}_i^*))$. Defining

$$w_i^* = \frac{\sum_{j=1}^2 \sum_{k=1}^{K_j} \lambda_{ijk}^* [z_{ijk}(w_{ijk} - \xi_{ijk})]}{\sum_{j=1}^2 \sum_{k=1}^{K_j} z_{ijk}^2 \lambda_{ijk}^*},$$

we have that w_i^* is sufficient for θ_i and $w_i^* | \theta_i \stackrel{\text{approx}}{\sim} N(\theta_i, 1/\sum_{j=1}^2 \sum_{k=1}^{K_j} z_{ijk}^2 \lambda_{ijk}^*)$. Further assume that the prior of δ is an m -state hidden Markov model (\mathcal{H}_m) with transition matrix \mathbf{C}_m , and $\theta_i | h_i = t \sim N(\nu_t, \kappa_t^2)$, for $t = 1, 2, \dots, m$, where h_i is the hidden state for θ_i . We marginalize over θ_i to obtain the approximate likelihood function

$$(3.2) \quad p(w_i^* | h_i = t) \stackrel{\text{approx}}{\sim} N\left(\nu_t, \kappa_t^2 + 1 / \sum_{j=1}^2 \sum_{k=1}^{K_j} z_{ijk}^2 \lambda_{ijk}^*\right)$$

where $t = 1, 2, \dots, m$.

The conditional prior probability, $P(h_i = t | h_j, j \neq i)$, for $t = 1, 2, 3$, can be computed from the transition matrix, \mathbf{C}_m , of the HMM \mathcal{H}_m .

The normalized product of the conditional prior probability and approximation (3.3) gives the approximate full conditional distribution of the differential state h_i , from which we propose a new value, $h_i^{(\text{prop})}$. We then propose a new value, $\theta_i^{(\text{prop})}$, from the approximate full conditional of θ_i given $h_i = h_i^{(\text{prop})}$. The proposed values $(h_i^{(\text{prop})}, \theta_i^{(\text{prop})})$ are jointly accepted or rejected by a Metropolis–Hastings probability [Gamerman and Lopes (2006)] to ensure that the post-burn-in MCMC samples represent draws from model posterior.

3.2. MCMC procedure. We iteratively generate MCMC samples of the chromosome-specific parameters by the following procedure:

1. The differential effects $\Delta_1, \dots, \Delta_I$ and latent differential states h_1, \dots, h_I are generated as in Section 3.1, given the expression magnitudes $\boldsymbol{\beta}$, subject-specific effects $\boldsymbol{\epsilon}$ (set to be 0 for nonpaired data) and treatment-replicate effects $\boldsymbol{\rho}$.

2. The mean expression magnitudes β_1, \dots, β_I and latent states s_1, \dots, s_I are generated as in Section 3.1, given the differential effects $\boldsymbol{\Delta}$, subject-specific effects $\boldsymbol{\epsilon}$ (set to be 0 for nonpaired data) and treatment-replicate effects $\boldsymbol{\rho}$.

3. For paired data, the subject-specific effects ε_{ik} for $i = 1, \dots, I$ and $k = 1, \dots, K$ are also generated by a similar Laplace approximation and Metropolis-Hastings procedure as in step 1.

4. Conditional on the mean expression magnitudes β_1, \dots, β_I , latent states s_1, \dots, s_I , and the fact that $\mu_2 - \mu_1 > \delta$, the hyperparameters μ_1 and μ_2 are jointly sampled from the restricted bivariate normal distribution using the R package *tmvt-norm* [Wilhelm and Manjunath (2013)].

5. For latent states $h = 1, 2, 3$, $s = 1, 2$, the hyperparameters $\phi_h, \tau_h^2, \sigma_s^2$ are all generated from their full conditional distributions by Gibbs sampling steps.

3.3. *Detection of DE genes.* For each chromosome, interest focuses on the latent vector of differential states h_1, \dots, h_I , where, as defined in equation (2.2), $h_i = 1$ ($h_i = 3$) represents an under-expressed (over-expressed) gene in treatment 2. We use the MCMC samples of the differential states to identify the DE genes while controlling for false discovery rate. Specifically, we use a Bayesian approach for controlling the FDR level first proposed by Newton et al. (2004), further studied by Müller, Parmigiani and Rice (2007), and subsequently applied in RNA-seq analysis by Lee et al. (2011).

Let q_0 be the desired nominal FDR level. In addition, let $r_i \in \{0, 1\}$ represent the unknown truth that gene i is differentially expressed ($r_i = 1$) or nondifferentially expressed ($r_i = 0$). Further, let p_i be the posterior probability that gene i is differentially expressed. Last, let $\delta_i \in \{0, 1\}$ denote the decision of calling gene i differentially expressed ($\delta_i = 1$) or nondifferentially expressed ($\delta_i = 0$). Using the MCMC output, we compute the estimate $\hat{p}_i = \Pr(r_i = 1 | \text{data})$ for genes $i = 1, \dots, N$ on all chromosomes.

A possible decision is to flag all genes with \hat{p}_i greater than or equal to a certain threshold p_0 . The resulting FDR would then be equal to

$$(3.3) \quad \text{FDR} = \frac{\sum_{i=1}^N (1 - r_i) 1_{(\hat{p}_i \geq p_0)}}{\sum_{i=1}^N 1_{(\hat{p}_i \geq p_0)}}.$$

Hence, the posterior expected FDR would be

$$(3.4) \quad \widehat{\text{FDR}} = \frac{\sum_{i=1}^N (1 - \hat{p}_i) 1_{(\hat{p}_i \geq p_0)}}{\sum_{i=1}^N 1_{(\hat{p}_i \geq p_0)}}.$$

Alternatively and more effectively than assigning a prespecified threshold p_0 , we may control the nominal FDR level q_0 [Müller, Parmigiani and Rice (2007), Newton et al. (2004)]. Specifically, first we rank genes in decreasing order of \hat{p}_i . Denote the ordered estimated posterior probabilities by $\hat{p}_{(1)} > \hat{p}_{(2)} > \dots > \hat{p}_{(N)}$. Thus, if we declare as differentially expressed the set of genes such that $\hat{p}_i \geq \hat{p}_{(d)}$, for each $d = 1, \dots, N$, then the corresponding posterior expected FDR will be

$$(3.5) \quad \widehat{\text{FDR}}_d = \frac{\sum_{i=1}^N (1 - \hat{p}_i) 1_{(\hat{p}_i \geq p_{(d)})}}{\sum_{i=1}^N 1_{(\hat{p}_i \geq p_{(d)})}} = \frac{\sum_{i=1}^d (1 - \hat{p}_{(i)})}{d}.$$

Finally, the decision rule for detecting DE genes is to flag all genes with $\widehat{\text{FDR}}_d < q_0$.

4. Simulation study. To compare the accuracy of `hmmSeq` with existing RNA-seq techniques, we performed two simulation studies. In the first study, the data were generated from a Poisson distribution. In the second simulation study, we generated the data from a negative binomial distribution. Three popular RNA-seq techniques were considered for comparisons: `edgeR` [Robinson, McCarthy and Smyth (2010)], `baySeq` [Hardcastle (2009)], and TSPM [Auer and Doerge (2011)]. The methods `edgeR` and `baySeq` have been implemented in R packages publicly available at <http://www.bioconductor.org>. R code for TSPM can be downloaded from <http://www.stat.purdue.edu/~doerge/software/TSPM.R>. The R code for `hmmSeq` is available in the Supplementary Materials [Cui et al. (2015)].

We first consider a simulation study with data generated from a Poisson distribution. For each of the following simulations, read counts were simulated for 12 chromosomes having 800 genes each, resulting in a total of 9600 genes. Six replicates of the set of read counts were generated for each of the two treatments. The replicate effects were assumed to be equal to the estimates for the biological replicates data of Zeng et al. (2012).

In the simulation study, the gene-specific magnitude factors β and the differential expression factors Δ were generated either from the hidden Markov model or finite mixture model with hyperparameters values given in Tables 1 and 2. The hyperparameters of the normal components were chosen to match the estimates for the biological replicates data of Zeng et al. (2012). The other hyperparameters were chosen according to our experience working with hidden Markov models [Guha, Li and Neuberg (2008)]. Let “F” denote a finite mixture model and “H” denote a hidden Markov model. We consider each model resulting from each possible combination of a FMM or a HMM for β and a FMM or a HMM for Δ in a total of 4 possible models. We denote each model with FF, FH, HF and HH, with the first letter indicating the process for β and the second letter indicating the process for Δ .

We perform model selection with the deviance information criterion (DIC) [Spiegelhalter et al. (2002)]. To evaluate the ability of DIC to discriminate among the four competing models, we have performed a simulation study. Specifically,

TABLE 1
Simulation study—parameters for generation of β

| | | Normal components | | | HMM transition matrix | FMM probability |
|---------|---------|-------------------|---------|--------------|---|-----------------|
| β | μ_1 | σ_1^2 | μ_2 | σ_2^2 | A $\begin{pmatrix} 0.50 & 0.50 \\ 0.05 & 0.95 \end{pmatrix}$ | P (0.1, 0.9) |
| | 1 | 0.37 | 3.91 | 2.4 | | |

TABLE 2
Simulation study—parameters for generation of Δ

| | | Normal components | | | | | HMM transition matrix | | | FMM probability |
|----------|----------|-------------------|------------|----------|------------|----------|-----------------------|------|--------------------|-----------------|
| Δ | ϕ_1 | τ_1^2 | τ_2^2 | ϕ_3 | τ_3^2 | B | | | Q | |
| | -0.4 | 0.013 | 0.01 | 0.4 | 0.013 | 0.50 | 0.25 | 0.25 | (0.22, 0.56, 0.22) | |
| | | | | | | 0.10 | 0.80 | 0.10 | | |
| | | | | | | 0.25 | 0.25 | 0.50 | | |

for each of the 4 possible models, we have simulated 30 data sets. After that, we have analyzed each simulated data set with the 4 different hmmSeq models: FF, FH, HF and HH. Then, for each simulated data set we have conducted DIC model selection. Table 3 presents for each true model the proportion of times that DIC has chosen each of the 4 competing models. As we can see from Table 3, the DIC chooses the correct model most of the time.

To compare hmmSeq with the other competing RNA-seq analysis methods, we consider their receiver operating characteristic (ROC) curves. The ROC curve of each method describes the relationship between the true positive rate (TPR) and the false positive rate (FPR) of gene detection. The TPR (which is also known as the sensitivity) is defined as the proportion of truly DE genes that are detected by the method. The FPR is defined as the proportion of non-DE genes that are erroneously identified as DE. The greater the area under the ROC curve, the greater the reliability of the method in detecting DE genes. For each simulation setup of FF, FH, HF and HH, we plot the ROC curves of (DIC picked) hmmSeq, edgeR, baySeq and TSPM averaged over 30 repetitions. Figure 3 displays the ROC curves for the methods hmmSeq (solid line), edgeR (dashed line), baySeq (dotted line) and TSPM (dot-dashed line) with the areas below the ROC curves being indicative of the relative accuracies of the methods in detecting DE genes. While edgeR beat the methods TSPM and baySeq in this simulation, hmmSeq achieves a substantially higher area under the ROC curve than the competing methods.

TABLE 3
Simulation study—performance of DIC-based model selection

| | | DIC chosen model | | | |
|------------|----|------------------|------|------|------|
| | | FF | FH | HF | HH |
| True model | FF | 0.73 | 0.10 | 0.07 | 0.10 |
| | FH | 0.00 | 0.77 | 0.00 | 0.23 |
| | HF | 0.07 | 0.03 | 0.80 | 0.10 |
| | HH | 0.00 | 0.10 | 0.00 | 0.90 |

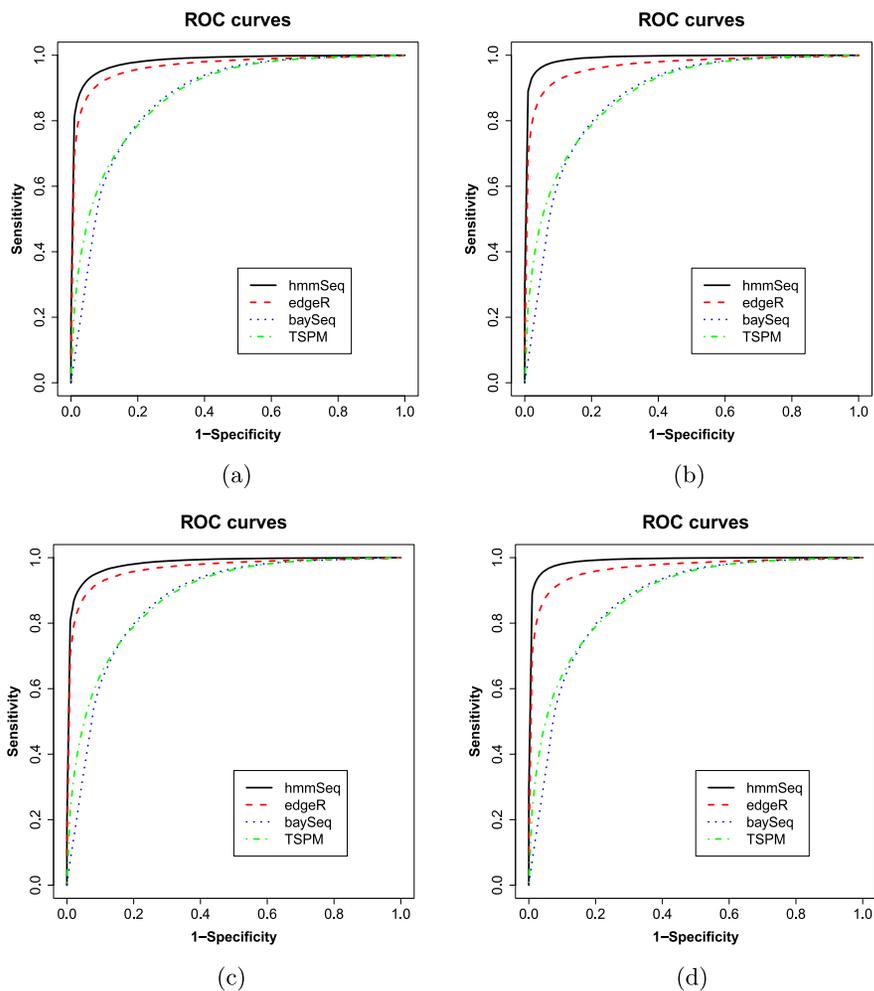


FIG. 3. For the simulation study of four setups FF (a), FH (b), HF (c) and HH (d), four panels depict the ROC curves for *hmmSeq*, *edgeR*, *baySeq* and *TSPM*. Results are averaged over 30 simulated data sets for each setup, where for each simulated data set we use the DIC-chosen *hmmSeq* model.

For each of the four competing methods, Figure 4 plots the observed FDR against the nominal FDR. Ideally, we would like to observe a 45 degree line through the origin in Figure 4 for each method. Observed FDR of *edgeR* is substantially smaller than the nominal FDR. The observed FDR for *TSPM* and *baySeq*, on the other hand, are quite liberal: The FDR for *TSPM* always exceeds 40%, while the FDR for *baySeq* exceeds 35% for most values of nominal FDR. Finally, FDR for *hmmSeq* is near and slightly lower than the 45 degree line. Therefore, *hmmSeq* is the method that performs best at controlling FDR.

To investigate the robustness of *hmmSeq* to overdispersed data, we simulated RNA-seq counts from a negative binomial distribution. This distribu-

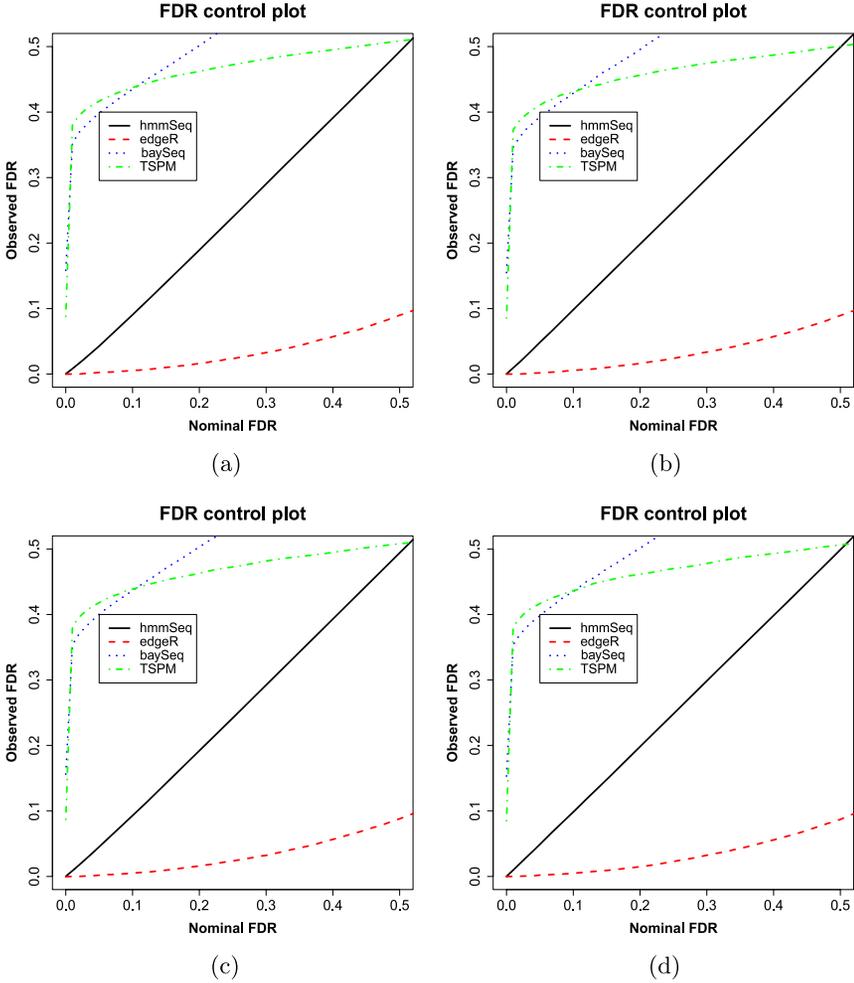


FIG. 4. For the simulation study, four panels depict the observed FDR versus nominal FDR for the methods *hmmSeq*, *TSPM*, *edgeR* and *baySeq* under four different simulation setups *FF* (a), *FH* (b), *HF* (c) and *HH* (d). Results are averaged over 30 simulated data sets for each setup, where for each simulated data set we use the DIC-chosen *hmmSeq* model. The proposed method controls the FDR closest and slightly lower to the 45 degree line.

tion is assumed by both *edgeR* and *baySeq*. Assume that $y|\lambda \sim \text{Poisson}(\lambda)$ and $\lambda|r, \psi \sim \text{gamma}(r, (1 - \psi)/\psi)$. Then, unconditionally, we obtain $y|r, \psi \sim \text{negative binomial}(r, \psi)$. The negative binomial mean is $m^{(1)} = r\psi/(1 - \psi)$ and the variance is $m^{(2)} = r\psi/(1 - \psi)^2$. The variance $m^{(2)} = m^{(1)}(1 + m^{(1)}/r)$ exceeds the mean $m^{(1)}$, reflecting overdispersion; $\zeta = 1/r$ is usually called the dispersion parameter.

The gene-specific magnitude factors β and the differential expression factors Δ were generated from a hidden Markov model with parameters values given in

Tables 1 and 2. To generate the negative binomial reads for each gene i , we first generated the mean $m_{ijkc}^{(1)} = \lambda_{ijkc}$ by equation (2.1) and the dispersion parameter ζ_i from a gamma distribution [as in, e.g., Kvam, Liu and Si (2012)]. To mimic the dispersion observed in real data sets, the shape parameter and scale parameter of the gamma distribution were estimated by the method of moments using the gene-wise dispersion estimates of Zeng et al. (2012) (biological replicates) available from edgeR. We computed the gamma distribution parameters $r = 1/\zeta$ and $\psi = \zeta\mu/(1 + \zeta\mu)$, and then hierarchically generated negative binomial read counts for 12 chromosomes having 800 genes each. This simulation procedure was replicated 30 times.

We fit four *hmmSeq* models to the negative binomial data. In addition, we fit edgeR, baySeq and TSPM models to the data. DIC chose the true model (HH in this case) 19 out of 30 times. The ROC curves and FDR controls are plotted in Figure 5. In the FDR control plot in Figure 5(b), we find that none of the methods are accurate. The *hmmSeq* FDR tends to be large for small nominal FDR, converging to the 45 degree line as the nominal level increases. In contrast, the observed FDR of baySeq and edgeR are mostly lower than the nominal FDR. For the ROC plot in Figure 5(a), *hmmSeq* achieves the highest area under the ROC curve than the competing methods, demonstrating its high reliability in detecting DE genes.

5. Data analysis. To illustrate the power and flexibility of our proposed RNA-seq analysis method, we have applied the *hmmSeq* method to analyze three data sets: Marioni et al. (2008) (technical replicates), Zeng et al. (2012) (biological

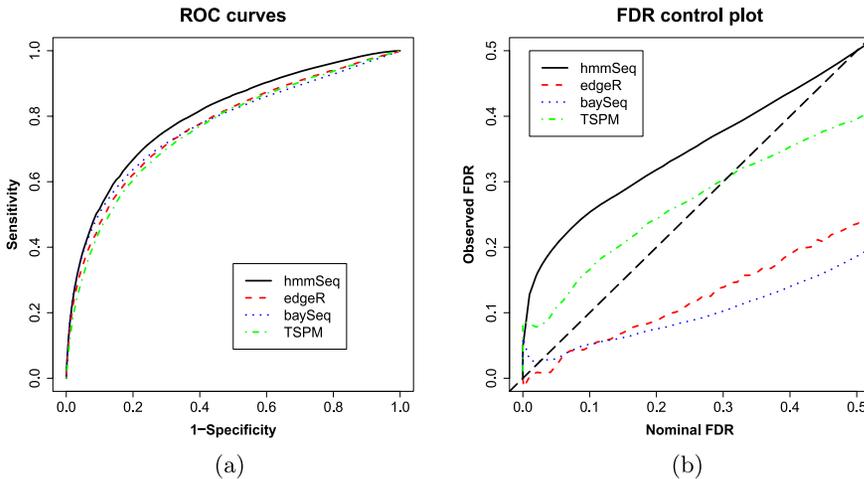


FIG. 5. For the negative binomial simulation study, the left panel depicts the ROC curves for DIC-selected *hmmSeq*, edgeR, baySeq and TSPM, and the right panel depicts the observed FDR versus nominal FDR for DIC-selected *hmmSeq*, TSPM, edgeR and baySeq. (a) ROC curves, (b) FDR control plot.

replicates) and Henn et al. (2013) (paired data). For each of these three data sets, the treatment-specific replicate effects ρ_{jk} are obtained by the upper-quartile normalizing technique of Bullard et al. (2010). In addition, we compare the results of the hmmSeq analysis with results of TSPM [Auer and Doerge (2011)], baySeq [Hardcastle (2009)] and edgeR [Robinson, McCarthy and Smyth (2010)] based on their publicly available R package implementations.

5.1. *Marioni et al. (2008) data set.* The Marioni et al. (2008) data set contains RNA-seq data for five technical replicates each of a single sample of kidney RNA (treatment 1) and liver RNA (treatment 2). Because genes with mostly small counts are not informative about differential expression, we have applied the filtering criterion of Auer and Doerge (2011) to eliminate the genes whose total read counts were less than 10. Additionally, the Y chromosome was ignored because many of its genes are transcribed on other chromosomes and the genders of the subjects are unknown. This yielded 17,076 genes for the analysis. Further, we applied the quantile normalization of Bullard et al. (2010) to preprocess the data. We have fitted four hmmSeq models (FF, FH, HF and HH) to the data. The DIC favors the FH model as the best, which implies neighboring genes dependence with respect to the differential expression parameter Δ . Thus, in the remainder of this section we present hmmSeq results based on the FH model.

We have applied the hmmSeq, edgeR, baySeq and TSPM methods to the Marioni et al. (2008) data set with a nominal FDR of $q_0 = 0.001$ [threshold adopted by Auer and Doerge (2011)]. Recall that the simulation study in Section 4 had indicated that the actual FDR of TSPM and baySeq is relatively insensitive to the value of q_0 and is substantially higher when q_0 is small. The sets of DE genes identified by the methods hmmSeq, edgeR, baySeq and TSPM are summarized in Figure 6(a). The method TSPM detected 9076 DE genes. In contrast, the hmmSeq method discovered only 2831 DE genes.

A closer examination sheds light on the differing sets of genes detected by TSPM and hmmSeq. Of the genes discovered by hmmSeq, as many as 2818 genes (99.5%) were also identified by the TSPM method. Focusing on the 6258 genes identified as DE by TSPM but not by hmmSeq, we find that TSPM flagged most of them as DE genes because they have extreme values of mean \log_2 -fold change (from output of TSPM) for the treatments. In particular, we have observed that for the 111 genes with \log_2 -fold change less than -20 , all five gene-specific reads for liver RNA were zero. Table 4 lists 10 randomly selected genes from this set, which reveals that the read counts for kidney, although positive, are not in hundreds or thousands like typical DE genes. Thus, TSPM tends to classify genes with 0 observations under any single condition as DE, while hmmSeq takes the variational magnitude into consideration. The result for edgeR lies between hmmSeq and TSPM.

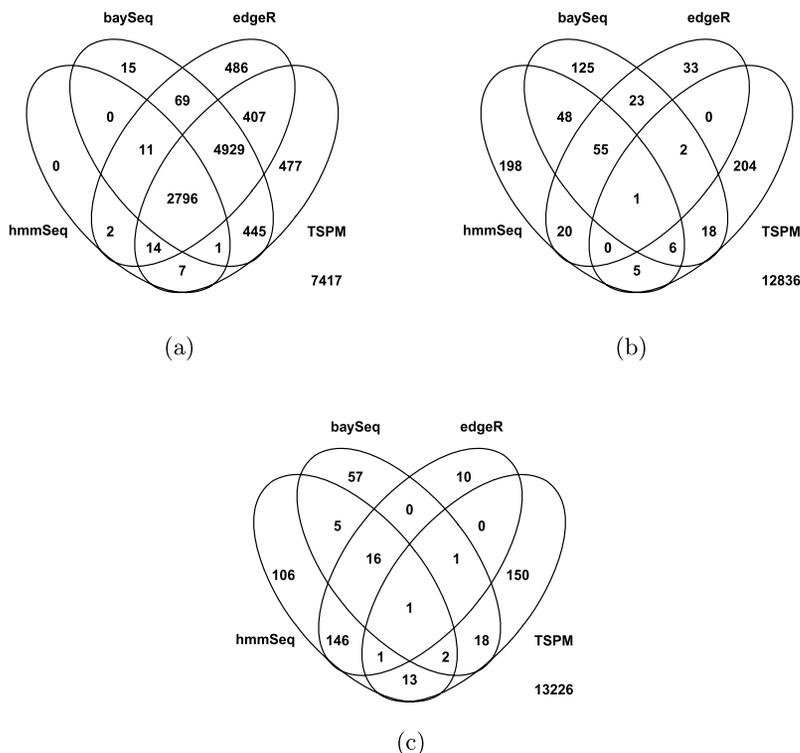


FIG. 6. Venn diagrams for the DE genes identified by the methods *hmmSeq*, *baySeq*, *edgeR* and *TSPM* in the data analyses. (a) Technical replicates of *Marioni et al. (2008)*, (b) biological replicates of *Zeng et al. (2012)*, (c) paired biological replicates of *Henn et al. (2013)*.

Similarly, all 33 genes with mean \log_2 -fold change greater than 20 have zero counts for kidney and relatively small counts for liver. The *hmmSeq* method called all of them non-DE, but *TSPM* classified them as DE genes due to their high mean \log_2 -fold changes. The former call seems more reasonable, given that the read counts of truly DE genes are typically several orders of magnitude higher.

5.2. *Zeng et al. (2012)* data set. We have applied *hmmSeq* to the data set of *Zeng et al. (2012)*, which contains samples from 2 regions of the human brain, frontal pole and hippocampus, with 6 biological replicates for each region. Data sets with biological replicates typically exhibit substantial *over-dispersion* or higher variability relative to the Poisson distribution. Researchers often model extra-Poisson variability using the binomial, negative binomial or Bayesian hierarchical Poisson models. We account for over-dispersion through hierarchical priors on the parameters in equation (2.1), for example, through random differential expression factors. There were 13,574 genes available for analysis after filtering (read sums for all the libraries did not exceed 9). Further, we have applied the

TABLE 4

Marioni et al. (2008) data set—ten randomly selected genes with mean \log_2 -fold change greater than 20 that were identified as DE genes by the TSPM method. R_jL_kK denotes j th run, k th replicate for kidney; R_jL_kL denotes j th run, k th replicate for liver

| Gene ID | R_1L_1K | R_1L_3K | R_1L_7K | R_2L_2K | R_2L_6K | R_1L_2L | R_1L_4L | R_1L_6L | R_1L_8L | R_2L_3L |
|-----------------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|
| ENSG00000198693 | 2 | 1 | 2 | 2 | 8 | 0 | 0 | 0 | 0 | 0 |
| ENSG00000162746 | 3 | 4 | 2 | 4 | 4 | 0 | 0 | 0 | 0 | 0 |
| ENSG00000168243 | 4 | 4 | 1 | 3 | 3 | 0 | 0 | 0 | 0 | 0 |
| ENSG00000188935 | 2 | 0 | 3 | 5 | 5 | 0 | 0 | 0 | 0 | 0 |
| ENSG00000173284 | 5 | 2 | 5 | 2 | 3 | 0 | 0 | 0 | 0 | 0 |
| ENSG00000114113 | 4 | 5 | 4 | 1 | 1 | 0 | 0 | 0 | 0 | 0 |
| ENSG00000169836 | 5 | 2 | 3 | 2 | 3 | 0 | 0 | 0 | 0 | 0 |
| ENSG00000170180 | 2 | 2 | 3 | 2 | 5 | 0 | 0 | 0 | 0 | 0 |
| ENSG00000186952 | 3 | 4 | 2 | 4 | 2 | 0 | 0 | 0 | 0 | 0 |
| ENSG00000164385 | 3 | 2 | 3 | 3 | 4 | 0 | 0 | 0 | 0 | 0 |

quantile normalization of Bullard et al. (2010) to preprocess the data. We have fitted four hmmSeq models (FF, FH, HF and HH) to the Zeng et al. (2012) data set. DIC chooses HH as the best model, which indicates neighboring genes dependence with respect both to the differential expression parameter Δ and to the expression level parameter β . Thus, in the remainder of this section we present hmmSeq results based on the HH model.

We have applied the hmmSeq, edgeR, baySeq and TSPM methods to the data set of Zeng et al. (2012) with a nominal FDR $q_0 = 0.05$. TSPM, edgeR and baySeq, respectively, identified 236, 134 and 278 DE genes. The hmmSeq technique identified 333 DE genes. The overlapping set of DE genes for the methods are summarized in Figure 6(b) and reveal a greater lack of agreement between the methods than for the Marioni et al. (2008) data set. Only 1 gene is identified as DE by all four methods. This low level of agreement is a result of the low overlap that TSPM has with the other methods. In contrast, hmmSeq has relatively large overlap both with edgeR (76 genes) and baySeq (110 genes).

We have investigated the biological implications of the results obtained with the hmmSeq analysis of differential expression of the hippocampus to the frontal pole. Though we expect a modest amount of differentially expressed genes, we do find some meaningful results that are supported in the literature. There is an increase in gene expression of Akt2 in the hippocampus compared to the frontal pole. Akt2 is a gene involved in insulin signaling, which occurs in the hippocampus [Agrawal and Gomez-Pinilla (2012), Robertson et al. (2010)]. In addition, Wnt7B is upregulated in the hippocampus where Wnt activity has been implicated in signaling of hippocampal synapses [Gogolla et al. (2009)]. Last, STAT5A is known to be expressed in the hippocampus [Kalita et al. (2013)], and our results show this upregulation. Taken all together, the results our hmmSeq method produced show biologically relevant genes when comparing the hippocampus to the frontal pole. In addition, we have conducted a functional analysis using DAVID [Huang, Sherman and Lempicki (2009a, 2009b)]; both our hmmSeq method and edgeR identified differentially expressed genes that are known to be expressed in brain tissue. The biological experiment presented tries to identify genes that are differentially expressed between two types of brain tissue. These results, taken in perspective with the biological experiment, suggest that the genes identified as differentially expressed via these two methods are relevant to the biological problem and help support the validity and accuracy of our predictions.

5.3. *Henn et al. (2013) paired data set.* Here we illustrate the application of the hmmSeq method to paired data sets with an analysis of a subset of an RNA-seq data set obtained by Henn et al. (2013) on immune response to a trivalent influenza vaccine. The original data set contains RNA-seq data from B cell samples for five subjects before vaccination (day 0) and for each of 10 days after vaccination (days 1 through 10). We consider a paired subset of three previously vaccinated subjects

in the original data set, where we compare gene expression before vaccination to gene expression after vaccination. Since peak B cell response usually appears 5–9 days post-vaccination, we apply our hmmSeq method to identify B cell gene differential expression between day 0 and day 7.

For the hmmSeq analysis, we first estimate the variance σ_ϵ^2 of subject-specific random effects from the data. Specifically, for each gene we have fitted a generalized linear mixed model with random subject effects, resulting in an estimated random effects variance for each gene. We then use the median of these estimates of random effects variances as an empirical Bayes estimate of σ_ϵ^2 . In addition, we have fitted the four hmmSeq models FF, FH, HF and HH. We have found that the DIC favors the FF model, that is, a finite mixture model without neighboring genes dependence. Thus, in the remainder of this section we present hmmSeq results based on the FF model.

We have analyzed this immune response data set using a nominal FDR of 0.05. To accommodate the paired data structure, in the edgeR analysis we include subject-specific fixed effects. Such edgeR analysis identifies 175 genes as differentially expressed. The TSPM that we used ignores the paired structure and treats all observations for each gene as independent, which identifies a total of 186 genes. Finally, in a paired baySeq analysis 100 genes are flagged to be DE. Figure 6(c) presents a Venn diagram that summarizes the results for TSPM, edgeR, baySeq and hmmSeq.

In order to further evaluate the competing methods, we compare their results to those found by Henn et al. (2013). Henn et al. (2013) used the RNA-seq data set from all 11 days, whereas we used only the data from days 0 and 7. Thus, here we use the results of Henn et al. (2013) as a benchmark. Specifically, Henn et al. (2013) identified a set of 742 genes as what they call the plasma cell gene signature (PCgs), that is, genes that have a common significant time-varying signature. Hence, in Table 5 we list the overlap of the PCgs set with the genes identified as differentially expressed by hmmSeq, edgeR, baySeq and TSPM. Our proposed hmmSeq method obtains the largest overlap with PCgs set (130 genes), and edgeR overlaps 121 genes with PCgs. We recall from Section 4 that hmmSeq and edgeR are the two methods with the highest area under the ROC curve. Thus, the overlap with the PCgs set shows the power of DE genes identification of the proposed hmmSeq method.

TABLE 5
Henn et al. (2013) data set—overlap of plasma cell gene signature (PCgs) set with genes identified by hmmSeq, edgeR, baySeq and TSPM

| | hmmSeq | edgeR | baySeq | TSPM |
|------|---------------|--------------|---------------|-------------|
| PCgs | 130 | 121 | 7 | 1 |

6. Conclusion. We propose *hmmSeq*, a method based on Bayesian hierarchical models for detecting DE genes between two treatments for paired or nonpaired data in RNA-seq analyses. The approach employs hidden Markov models to account for the statistical dependence between the gene counts of neighboring genes observed in many RNA-seq data sets. The *hmmSeq* model can be applied to studies with either biological or technical replicates, automatically adjusting for any overdispersion relative to the Poisson distribution. Through simulated and real data sets, we compare and contrast the performance of *hmmSeq* with some well-known methods in the literature, demonstrating the reliability and success of our approach in the identification of DE genes.

We have developed DIC-based model selection to decide for each data set whether HMM or FMM should be used to model gene expression magnitude and/or differential expression. For the [Marioni et al. \(2008\)](#) data set the DIC-selected model is the FH model, for the [Zeng et al. \(2012\)](#) data set the DIC-selected model is the HH model, and for the [Henn et al. \(2013\)](#) data set the DIC-selected model is the FF model. Thus, for one data set there appears to be neighboring genes dependence in expression magnitude. Even more important, for two data sets there is evidence of neighboring genes dependence in differential expression. This co-differential expression can be justified by how ancient species organized their genomes and by evolution. Specifically, more ancient species, such as bacteria, organize their genomes based on operons, where genes involved in the same process or needed at the same time are transcribed in tandem [[Alberts et al. \(1994\)](#)]. Throughout evolution, operons have been divided into individual genes, but genes involved in the same process now reside in gene clusters [[Hurst, Pál and Lercher \(2004\)](#)]. Thus, neighboring genes tend to be jointly differentially expressed.

To further examine spatial genomic dependence (and clustering) among detected DE genes, we have devised the following statistical test. Consider any detected DE gene and the next detected DE gene in the chromosome as neighboring DE genes. Consider the distance between two neighboring DE genes as the number of non-DE genes between them. If there is no spatial dependence, then all the distances between any two neighboring DE genes should be a random sample from a geometric distribution. Hence, to test for spatial dependence, we collect all the distances between neighboring DE genes and conduct a goodness-of-fit test of the hypothesis that the empirical distribution equals the null theoretical geometric distribution. We use this procedure to test for spatial genomic dependence for DE calls from *edgeR* and *hmmSeq*. First, we performed this test for the [Henn et al. \(2013\)](#) data set for which *hmmSeq* prefers finite mixture model and spatial independence. The spatial genomic dependence test for DE calls from *edgeR* and *hmmSeq* yields p -values equal to 0.2201 and 0.5178, respectively, further supporting *hmmSeq* suggestion of spatial independence. Second, we performed this test for the two real data sets for which *hmmSeq* prefers spatial dependence, that is, the [Marioni et al. \(2008\)](#) and the [Zeng et al. \(2012\)](#) data sets. For the [Marioni et al. \(2008\)](#) data set, the spatial dependence test for DE calls from both *hmmSeq* and

edgeR yield p -values smaller than $2.2e-16$. That is, even though edgeR does not account for spatial dependence, its detected DE genes for the [Marioni et al. \(2008\)](#) data set cluster spatially. For the [Zeng et al. \(2012\)](#) data set, edgeR only detected 134 DE genes which did not provide enough power for the goodness-of-fit data set (p -value greater than 0.9). In contrast, there is strong statistical evidence that the 333 genes identified by `hmmSeq` as DE cluster spatially (p -value smaller than $2.2e-16$). Therefore, these data sets point to the need to consider spatial genomic dependence in studies of differential gene expression.

In addition to genomic spatial dependence among genes based on genomic position, for future research work we plan to extend `hmmSeq` to include other sources of dependence among genes. Recent experimental techniques such as HiC and ChIA-PET allow for the identification of explicit promoter–promoter and promoter–enhancer–promoter interactions [[Edelman and Fraser \(2012\)](#), [Mercer and Mattick \(2013\)](#), [van Arensbergen, van Steensel and Bussemaker \(2014\)](#)]. In addition, we note that genes that belong to the same active functional pathways tend to be co-expressed [[Tegge, Caldwell and Xu \(2012\)](#)]. This extension of `hmmSeq` may need a non-Markovian spatial correlation model. We leave this challenging inferential problem to future research.

DE gene call lists are frequently used in downstream pathway function calls in what is known as functional enrichment analysis. Because functional enrichment analysis methods usually assume independence of DE gene calls, caution needs to be taken when using the DE gene call lists generated by `hmmSeq`. When `hmmSeq` decides that the best model is a finite mixture model without spatial dependence, then one can use `hmmSeq`'s DE gene call list without any concern. However, when `hmmSeq` decides that a spatial dependence model is warranted, then the assumption of independence no longer holds. This opens up a tremendous opportunity for future research that performs joint differential expression gene calls and functional enrichment analysis. We envision this joint analysis may be implemented by extending `hmmSeq` to incorporate information on functional pathway networks.

In terms of computational time, on a desktop with a 2.3 GHz processor and 4 GB memory, `hmmSeq` takes approximately 3 hours to analyze a 1200-gene chromosome. Although it does take a longer time than other methods, `hmmSeq` often achieves a higher accuracy of DE gene detection than other methods by a realistic model that allows for spatial genomic dependence. Moreover, the computational time of all the considered statistical methods is negligible when compared to the time (on the order of months or years) required by subject-matter scientists to perform experiments to obtain RNA-seq data. Furthermore, to limit the computational time the `hmmSeq` analysis can be performed in parallel for individual chromosomes. Finally, when compared to the high costs of RNA-seq extraction, the information gains obtained by the `hmmSeq` methodology seem well worth the relatively low computational costs.

The `hmmSeq` method we propose relies on a single user-specified “tuning” parameter q_0 , that is, the nominal false discovery rate. A default value for q_0 between 0.001 or 0.05 has produced satisfactory results for all the data sets, real or simu-

lated, that we have analyzed, facilitating “black box” applications of `hmmSeq`. Future work will focus on extending `hmmSeq` to investigations with more than two treatments. An R package implementing the `hmmSeq` framework will be submitted to CRAN upon publication of the manuscript.

Acknowledgments. The authors would like to thank Martin Zand and Stephen Welle at the University of Rochester Medical Center for providing the vaccine data in Section 5.3.

SUPPLEMENTARY MATERIAL

Supplement to “`hmmSeq`: A hidden Markov model for detecting differentially expressed genes from RNA-seq data.” (DOI: [10.1214/15-AOAS815SUPP](https://doi.org/10.1214/15-AOAS815SUPP); .zip). The R code for `hmmSeq`.

REFERENCES

- AGRAWAL, R. and GOMEZ-PINILLA, F. (2012). ‘Metabolic syndrome’ in the brain: Deficiency in omega-3 fatty acid exacerbates dysfunctions in insulin receptor signalling and cognition. *J. Gen. Physiol.* **590** 2485–2499.
- ALBERTS, B., BRAY, D., LEWIS, J., RAFF, M., ROBERTS, K. and WATSON, J. D. (1994). *Molecular Biology of the Cell*. Garland Science, New York.
- AUER, P. L. and DOERGE, R. W. (2011). A two-stage Poisson model for testing RNA-Seq data. *Stat. Appl. Genet. Mol. Biol.* **10** Art. 26, 28. [MR2800690](https://doi.org/10.1214/11-AM/12800690)
- AUER, P. L., SRIVASTAVA, S. and DOERGE, R. W. (2012). Differential expression—the next generation and beyond. *Brief. Funct. Genomics* **11** 57–62.
- BENJAMINI, Y. and HOCHBERG, Y. (1995). Controlling the false discovery rate: A practical and powerful approach to multiple testing. *J. Roy. Statist. Soc. Ser. B* **57** 289–300. [MR1325392](https://doi.org/10.1111/j.1467-9868.1995.tb01704.x)
- BLEKHMEN, R., MARIONI, J. C., ZUMBO, P., STEPHENS, M. and GILAD, Y. (2010). Sex-specific and lineage-specific alternative splicing in primates. *Genome Res.* **20** 180–189.
- BULLARD, J. H., PURDOM, E., HANSEN, K. D. and DUDOIT, S. (2010). Evaluation of statistical methods for normalization and differential expression in mRNA-seq experiments. *BMC Bioinformatics* **11** 94.
- CARON, H., VAN SCHAİK, B., VAN DER MEE, M., BAAS, F., RIGGINS, G., VAN SLUIS, P., HERMUS, M.-C., VAN ASPEREN, R., BOON, K., VOUTE, P. A. et al. (2001). The human transcriptome map: Clustering of highly expressed genes in chromosomal domains. *Science* **291** 1289–1292.
- CHIB, S. and GREENBERG, E. (1994). Bayes inference in regression models with ARMA(p, q) errors. *J. Econometrics* **64** 183–206. [MR1310523](https://doi.org/10.1016/0927-6460(94)90052-3)
- CUI, S., GUHA, S., FERREIRA, M. and TEGGE, A. N. (2015). Supplement to “`hmmSeq`: A hidden Markov model for detecting differentially expressed genes from RNA-seq data.” DOI:[10.1214/15-AOAS815SUPP](https://doi.org/10.1214/15-AOAS815SUPP).
- EDELMAN, L. B. and FRASER, P. (2012). Transcription factories: Genetic programming in three dimensions. *Curr. Opin. Genet. Dev.* **22** 110–114.
- FRÜHWIRTH-SCHNATTER, S. (2006). *Finite Mixture and Markov Switching Models*. Springer, New York. [MR2265601](https://doi.org/10.1007/978-1-4020-5111-1)
- GAMERMAN, D. and LOPES, H. F. (2006). *Markov Chain Monte Carlo: Stochastic Simulation for Bayesian Inference*, 2nd ed. Chapman & Hall/CRC, Boca Raton, FL. [MR2260716](https://doi.org/10.1080/00036810600571616)
- GOGOLLA, N., GALIMBERTI, I., DEGUCCI, Y. and CARONI, P. (2009). Wnt signaling mediates experience-related regulation of synapse numbers and mossy fiber connectivities in the adult hippocampus. *Neuron* **62** 510–525.

- GUHA, S., LI, Y. and NEUBERG, D. (2008). Bayesian hidden Markov modeling of array CGH data. *J. Amer. Statist. Assoc.* **103** 485–497. [MR2523987](#)
- HARDCASTLE, T. J. (2009). baySeq: Empirical Bayesian analysis of patterns of differential expression in count data. R package version 1.10.0.
- HARDCASTLE, T. J. and KELLY, K. A. (2010). baySeq: Empirical Bayesian methods for identifying differential expression in sequence count data. *BMC Bioinformatics* **11** 422.
- HENN, A. D., WU, S., QIU, X., RUDA, M., STOVER, M., YANG, H., LIU, Z., WELLE, S. L., HOLDEN-WILTSE, J., WU, H. and ZAND, M. S. (2013). High-resolution temporal response patterns to influenza vaccine reveal a distinct human plasma cell gene signature. *Sci. Rep.* **3** 2327.
- HUANG, D. W., SHERMAN, B. T. and LEMPICKI, R. A. (2009a). Bioinformatics enrichment tools: Paths toward the comprehensive functional analysis of large gene lists. *Nucleic Acids Res.* **37** 1–13.
- HUANG, D. W., SHERMAN, B. T. and LEMPICKI, R. A. (2009b). Systematic and integrative analysis of large gene lists using DAVID bioinformatics resources. *Nat. Protoc.* **4** 44–57.
- HURST, L. D., PÁL, C. and LERCHER, M. J. (2004). The evolutionary dynamics of eukaryotic gene order. *Nat. Rev. Genet.* **5** 299–310.
- KALITA, A., GUPTA, S., SINGH, P., SUROLIA, A. and BANERJEE, K. (2013). IGF-1 stimulated upregulation of cyclin D1 is mediated via STAT5 signaling pathway in neuronal cells. *IUBMB Life* **65** 462–471.
- KARLEBACH, G. and SHAMIR, R. (2008). Modelling and analysis of gene regulatory networks. *Nat. Rev., Mol. Cell Biol.* **9** 770–780.
- KVAM, V. M., LIU, P. and SI, Y. (2012). A comparison of statistical methods for detecting differentially expressed genes from RNA-seq data. *Am. J. Bot.* **99** 248–256.
- LANGMEAD, B., HANSEN, K. D., LEEK, J. T. et al. (2010). Cloud-scale RNA-sequencing differential expression analysis with Myrna. *Genome Biol.* **11** R83.
- LEE, J., JI, Y., LIANG, S., CAI, G. and MÜLLER, P. (2011). On differential gene expression using RNA-seq data. *Cancer Inform.* **10** 205–215.
- LOUHIMO, R. and HAUTANIEMI, S. (2011). CNAmets: An R package for integrating copy number, methylation and expression data. *Bioinformatics* **27** 887–888.
- MACDONALD, I. L. and ZUCCHINI, W. (1997). *Hidden Markov and Other Models for Discrete-Valued Time Series. Monographs on Statistics and Applied Probability* **70**. Chapman & Hall, London. [MR1692202](#)
- MARIONI, J. C., MASON, C. E., MANE, S. M., STEPHENS, M. and GILAD, Y. (2008). RNA-seq: An assessment of technical reproducibility and comparison with gene expression arrays. *Genome Res.* **18** 1509–1517.
- MERCER, T. R. and MATTICK, J. S. (2013). Understanding the regulatory and transcriptional complexity of the genome through structure. *Genome Res.* **23** 1081–1088.
- MICHALAK, P. (2008). Coexpression, coregulation, and cofunctionality of neighboring genes in eukaryotic genomes. *Genomics* **91** 243–248.
- MÜLLER, P., PARMIGIANI, G. and RICE, K. (2007). FDR and Bayesian multiple comparisons rules. In *Bayesian Statistics* **8** (J. M. Bernardo, S. Bayarri, J. O. Berger, A. Dawid, D. Heckerman, A. F. M. Smith and M. West, eds.) 349–370. Oxford Univ. Press, Oxford. [MR2433200](#)
- NEWTON, M. A., NOUEIRY, A., SARKAR, D. and AHLQUIST, P. (2004). Detecting differential gene expression with a semiparametric hierarchical mixture method. *Biostatistics* **5** 155–176.
- PE’ER, D. and HACOEN, N. (2011). Principles and strategies for developing network models in cancer. *Cell* **144** 864–873.
- RABINER, L. R. (1989). A tutorial on hidden Markov models and selected applications in speech recognition. *Proc. IEEE* **77** 257–286.
- ROBERTSON, S. D., MATTHIES, H. J., OWENS, W. A., SATHANANTHAN, V., CHRISTIANSON, N. S. B., KENNEDY, J. P., LINDSLEY, C. W., DAWS, L. C. and GALLI, A. (2010).

- Insulin reveals akt signaling as a novel regulator of norepinephrine transporter trafficking and norepinephrine homeostasis. *J. Neurosci.* **30** 11305–11316.
- ROBINSON, M. D., MCCARTHY, D. J. and SMYTH, G. K. (2010). edgeR: A bioconductor package for differential expression analysis of digital gene expression data. *Bioinformatics* **26** 139–140.
- ROBINSON, M. D. and SMYTH, G. K. (2007). Moderated statistical tests for assessing differences in tag abundance. *Bioinformatics* **23** 2881–2887.
- ROBINSON, M. D. and SMYTH, G. K. (2008). Small-sample estimation of negative binomial dispersion, with applications to SAGE data. *Biostatistics* **9** 321–332.
- SCOTT, S. L. (2002). Bayesian methods for hidden Markov models. *J. Amer. Statist. Assoc.* **97** 337–351.
- SI, Y. and LIU, P. (2013). An optimal test with maximum average power while controlling FDR with application to RNA-Seq data. *Biometrics* **69** 594–605. [MR3106587](#)
- SINGER, G. A., LLOYD, A. T., HUMNIECKI, L. B. and WOLFE, K. H. (2005). Clusters of co-expressed genes in mammalian genomes are conserved by natural selection. *Mol. Biol. Evol.* **22** 767–775.
- SPIEGELHALTER, D. J., BEST, N. G., CARLIN, B. P. and VAN DER LINDE, A. (2002). Bayesian measures of model complexity and fit. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **64** 583–639. [MR1979380](#)
- STOREY, J. D. and TIBSHIRANI, R. (2003). Statistical significance for genomewide studies. *Proc. Natl. Acad. Sci. USA* **100** 9440–9445. [MR1994856](#)
- TEGGE, A. N., CALDWELL, C. W. and XU, D. (2012). Pathway correlation profile of gene–gene co-expression for identifying pathway perturbation. *PLoS ONE* **7** e52127.
- TITTERINGTON, D. M., SMITH, A. F. M. and MAKOV, U. E. (1985). *Statistical Analysis of Finite Mixture Distributions*. Wiley, Chichester. [MR0838090](#)
- VAN ARENSBERGEN, J., VAN STEENSEL, B. and BUSSEMAKER, H. J. (2014). In search of the determinants of enhancer–promoter interaction specificity. *Trends Cell Biol.* **24** 695–702.
- WILHELM, S. and MANJUNATH, B. G. (2013). tmvtnorm: Truncated multivariate normal and student t distribution. R package version 1.4-8.
- ZEGER, S. L. and KARIM, M. R. (1991). Generalized linear models with random effects; a Gibbs sampling approach. *J. Amer. Statist. Assoc.* **86** 79–86. [MR1137101](#)
- ZENG, J., KONOPKA, G., HUNT, B. G., PREUSS, T. M., GESCHWIND, D. and YI, S. V. (2012). Divergent whole-genome methylation maps of human and chimpanzee brains reveal epigenetic basis of human regulatory evolution. *Am. J. Hum. Genet.* **91** 455–465.
- ZHAO, S., FUNG-LEUNG, W.-P., BITTNER, A., NGO, K. and LIU, X. (2014). Comparison of RNA-seq and microarray in transcriptome profiling of activated t cells. *PLoS ONE* **9** e78644.

S. CUI
S. GUHA
DEPARTMENT OF STATISTICS
UNIVERSITY OF MISSOURI
146 MIDDLEBUSH HALL
COLUMBIA, MISSOURI 65211-6100
USA
E-MAIL: scrr2@mail.missouri.edu
guhasu@missouri.edu

M. A. R. FERREIRA
DEPARTMENT OF STATISTICS
VIRGINIA TECH
BLACKSBURG, VIRGINIA 24061
USA
E-MAIL: marf@vt.edu

A. N. TEGGE
COMPUTER SCIENCE
VIRGINIA TECH
BLACKSBURG, VIRGINIA 24061
USA
E-MAIL: ategge@vt.edu