

IsoDOT Detects Differential RNA-Isoform Expression/Usage With Respect to a Categorical or Continuous Covariate With High Sensitivity and Specificity

Wei SUN, Yufeng LIU, James J. CROWLEY, Ting-Huei CHEN, Hua ZHOU, Haitao CHU, Shunping HUANG, Pei-Fen KUAN, Yuan LI, Darla MILLER, Ginger SHAW, Yichao WU, Vasyi ZHABOTYNSKY, Leonard MCMILLAN, Fei ZOU, Patrick F. SULLIVAN, and Fernando Pardo-Manuel DE VILLENA

We have developed a statistical method named IsoDOT to assess differential isoform expression (DIE) and differential isoform usage (DIU) using RNA-seq data. Here isoform usage refers to relative isoform expression given the total expression of the corresponding gene. IsoDOT performs two tasks that cannot be accomplished by existing methods: to test DIE/DIU with respect to a continuous covariate, and to test DIE/DIU for one case versus one control. The latter task is not an uncommon situation in practice, for example, comparing the paternal and maternal alleles of one individual or comparing tumor and normal samples of one cancer patient. Simulation studies demonstrate the high sensitivity and specificity of IsoDOT. We apply IsoDOT to study the effects of haloperidol treatment on the mouse transcriptome and identify a group of genes whose isoform usages respond to haloperidol treatment. Supplementary materials for this article are available online.

KEY WORDS: Differential isoform expression; Differential isoform usage; Isoform; Penalized regression; RNA-seq

In the genomes of higher eukaryotes, the DNA sequence of a gene often includes multiple exons that are separated by introns. A multi-exon gene may encode several RNA isoforms, each with a unique subset of exons. Recent studies have shown that more than 90% of human genes have multiple RNA isoforms which may be differentially expressed across tissues or developmental stages (Pan et al. 2008; Wang et al. 2008), and about 75% of human genes produce multiple RNA isoforms within a given cell type (Djebali et al. 2012). Therefore, study of RNA-isoform expression and its regulation is of great importance to understand the functional complexity of a living organism, the evolutionary changes in transcriptome (Barbosa-Morais et al. 2012), and the genomic basis of human diseases (Wang and Cooper 2007).

Gene expression has traditionally been measured by microarrays or exon arrays, most of which provide just a handful of probes per gene and poor resolution to distinguish multiple iso-

forms. (Purdom et al. 2008; Richard et al. 2010). RNA sequencing (RNA-seq), on the other hand, provides much better data for this purpose (Wang, Gerstein, and Snyder 2009). In RNA-seq, fragments of RNA molecules (typically 200–500 bps long) are reverse transcribed and amplified, followed by sequencing of one or both ends (single vs paired-end). Each sequence is referred to as a read, which could be 30–150 bps or even longer. Reads are then mapped to a reference genome and the number of fragments overlapping each gene is counted. The expression of the j th gene in the i th sample is measured by a normalized fragment count after adjusting for read-depth of the i th sample and the length of the j th gene (Mortazavi et al. 2008).

The primary challenge with studying RNA isoforms is that we cannot directly observe the expression of each RNA isoform. More specifically, an RNA-seq fragment may be compatible with more than one RNA isoform, and thus, we cannot unambiguously assign it to an RNA isoform. Several methods have been developed to address this challenge (Xing et al. 2006; Jiang and Wong 2009; Richard et al. 2010; Trapnell et al. 2010; Li et al. 2010; Katz et al. 2010; Pachter 2011; Roberts et al. 2011; Salzman, Jiang, and Wong 2011; Chen 2012). Moreover, the annotation of RNA isoforms may not be complete or accurate and thus, one may need to reconstruct transcriptome annotation using RNA-seq data (Denoëud et al. 2008). Simultaneous transcriptome reconstruction and isoform abundance estimation can be achieved using different approaches, including penalized regression methods (Bohnert and Rättsch 2010; Xia et al. 2011; Li, Feng, and Jiang 2011b; Li et al. 2011a), where each possible isoform is treated as a covariate in a regression problem. Interested readers are referred to Alamancos, Agirre, and Eyras (2014)

Wei Sun is corresponding author (E-mail: wsun@bios.unc.edu), Ting-huei Chen (E-mail: tchen@live.unc.edu), Pei-Fen Kuan (E-mail: pfkuan@email.unc.edu), and Fei Zou (E-mail: fzou@bios.unc.edu), Department of Biostatistics, University of North Carolina, Chapel Hill, NC 27599. Yufeng Liu is Professor, Lineberger Comprehensive Cancer Center, University of North Carolina, 354 Hanes Hall, Chapel Hill, NC 27599 (E-mail: yfliu@email.unc.edu). James Crowley (E-mail: crowley@unc.edu), Darla Miller (E-mail: darla_miller@med.unc.edu), Ginger Shaw (E-mail: shawg@email.unc.edu), Vasyi Zhabotynsky (E-mail: vzhaboty@bios.unc.edu), Patrick Sullivan (E-mail: psullivan@med.unc.edu), and Fernando Pardo-Manuel de Villena (E-mail: fernando@med.unc.edu), Department of Genetics, University of North Carolina, Chapel Hill, NC 27599. Hua Zhou (E-mail: hzhou3@ncsu.edu), Yuan Li (E-mail: yli16@ncsu.edu), and Yichao Wu (E-mail: wu@stat.ncsu.edu), Department of Statistics, North Carolina State University, Raleigh, NC 27695. Haitao Chu, Department of Biostatistics, University of Minnesota, Minneapolis, MN 55455 (E-mail: chux0051@umn.edu). Shunping Huang (E-mail: mark.sphuang@gmail.com) and Leonard McMillan (E-mail: mcmillan@unc.edu), Department of Computer Science, University of North Carolina, Chapel Hill, NC 27599. This work was partially supported by NIH grants GM105785, CA167684, CA149569, MH101819, GM074175, P50 HG006582, and K01MH094406.

Color versions of one or more of the figures in the article can be found online at www.tandfonline.com/r/jasa.

for a comprehensive list of relevant statistical/computational methods.

Differential isoform expression (DIE) and differential isoform usage (DIU) are related but distinct concepts. DIE assesses the difference of absolute expression in isoform level. In contrast, DIU assesses the difference of relative expression in isoform level. For example, if the expression of two isoforms of one gene are 10 and 20 in control and 50 and 100 in case, then there is DIE but no DIU because the relative expression of the first isoform is $1/3$ in both case and control. Although many methods have been developed to estimate RNA isoform expression, only a few methods have been developed to assess DIE or DIU while modeling biological variability and accounting for the uncertainty of isoform expression estimation. These methods include BitSeq (Glaus, Honkela, and Rattray 2012), Cuffdiff2 (Trapnell et al. 2013), and EBseq (Leng et al. 2013). All three methods are designed for two-group or multi-group comparisons with multiple samples per group. BitSeq (Bayesian inference of transcripts from sequencing data) adopts a two-stage approach. The first stage is isoform expression estimation within each sample using a Bayesian MCMC method. The second stage is to assess differential expression of each isoform using the posterior samples from the first stage. Cuffdiff2 employs a likelihood-based approach for isoform expression estimation and relevant hypothesis testing. For each gene, Cuffdiff2 first estimates expectation and covariance of the expression of multiple isoforms, and then uses these estimates to assess differential expression of this gene or DIU of its isoforms. For differential expression, Cuffdiff2 constructs a test statistic of log fold change, standardized by its standard error. Cuffdiff2 offers two tests for DIU: one for all the isoforms sharing a transcription starting site (TSS) and one for differential usage of TSSs. The test statistic for DIU is the square root of the Jensen–Shannon divergence, divided by its standard error. While both BitSeq and Cuffdiff2 first estimate isoform expression and then perform hypothesis testing, EBSeq uses isoform expression estimates from other methods. For two-group or multi-group comparisons, EBSeq assumes the (rounded) isoform expression estimate follows a negative binomial distribution with group-specific mean and overdispersion. EBSeq stratifies all isoforms into multiple categories to allow category-specific mean-variance relations. These isoform categories are constructed based on the difficulty of isoform expression estimation. For example, genes with one, two, or more isoforms may form three categories.

BitSeq, Cuffdiff2, and EBseq all address an important issue for DIE: to account for the uncertainty inherent in the isoform expression estimation process. However, there are two types of commonly encountered tasks that cannot be accomplished by these methods: to assess DIE with respect to a continuous variable, for example, age or additive coding of genotype (i.e., 0, 1, 2, for genotype AA, AB, and BB), and to assess DIE across two groups with only one sample per group, which is not an uncommon situation in real data studies. For example, one may wish to compare isoform expression between the paternal and maternal alleles of an individual or between normal and cancer tissues of a patient. In such situations, the RNA-seq data allow a valid statistical test, although the population for statistical inference is limited to the tested case and control (i.e., what happens if we collect more RNA-seq fragments from this

case and this control) rather than the general case and control populations (i.e., what happens if we collect more samples from case or control population). BitSeq and EBseq cannot compare two groups with one case and one control. Cuffdiff2 provides an ad-hoc implementation for this problem. Specifically, when there is one case and one control, Cuffdiff2 estimates isoform expression variance by combining case and control, which implicitly assumes most isoforms are not differentially expressed. Therefore, it is expected that Cuffdiff2 would yield conservative p -values and limited power in this situation, which is confirmed in our simulation studies.

In this article, we develop a statistical method named IsoDOT, which assesses DIE or DIU using RNA-seq data and addresses the aforementioned two tasks that cannot be accomplished by existing methods. IsoDOT treats all the RNA isoforms of a gene (or a transcript cluster of a few overlapping genes) as a unit and tests whether any of these RNA isoforms is associated with a covariate of interest. Alternative strategies would be to assess differential expression or differential usage of each exon set or each RNA isoform. For testing at the exon set level, the number of tests is much larger than gene-level testing, which increases the burden on multiple testing correction. In fact, multiple testing correction is also more challenging because multiple exon sets of the same gene often have correlated expression. For isoform-level testing, the major challenge is to incorporate the uncertainty in isoform expression estimation into the testing step. It is possible that two isoforms are very similar and thus, available data cannot distinguish them. Therefore, differential expression testing for these two isoforms separately is problematic. By performing testing per transcript cluster, IsoDOT bypasses the limitation of exon-set-level or RNA-isoform-level testing. After transcript clusters with significant DIE or DIU are identified, one may follow up on these transcript clusters to identify differentially expressed exon sets (Anders, Reyes, and Huber 2012) or isoforms (Glaus, Honkela, and Rattray 2012; Leng et al. 2013; Trapnell et al. 2013).

MATERIALS AND METHODS

An Overview

We assume that the locations and sizes of all the exons of a gene are known. If needed, one can use existing software (e.g., TopHat; Trapnell, Pachter, and Salzberg 2009) to detect previously unknown exons. The inputs of our method are the bam files of all samples. From each bam file, we derive the number of RNA-seq fragments overlapping each exon set (an exon set includes one or more adjacent or non adjacent exons) and the distribution of RNA-seq fragments' lengths (Figure 1(a)). IsoDOT estimates RNA isoform expression across all samples, and outputs two p -values per gene: one for testing DIE and one for testing DIU. The DIE test asks whether the absolute expression of any isoform of a gene is associated with the covariate of interest. In contrast, the DIU test asks, after adjusting for total expression of the corresponding gene, whether the relative expression of any isoform of this gene is associated with the covariate of interest.

As part of IsoDOT, we have developed a penalized regression method, named IsoDetector, to estimate RNA isoform expression. In contrast to existing methods (Bohnert and Rättsch

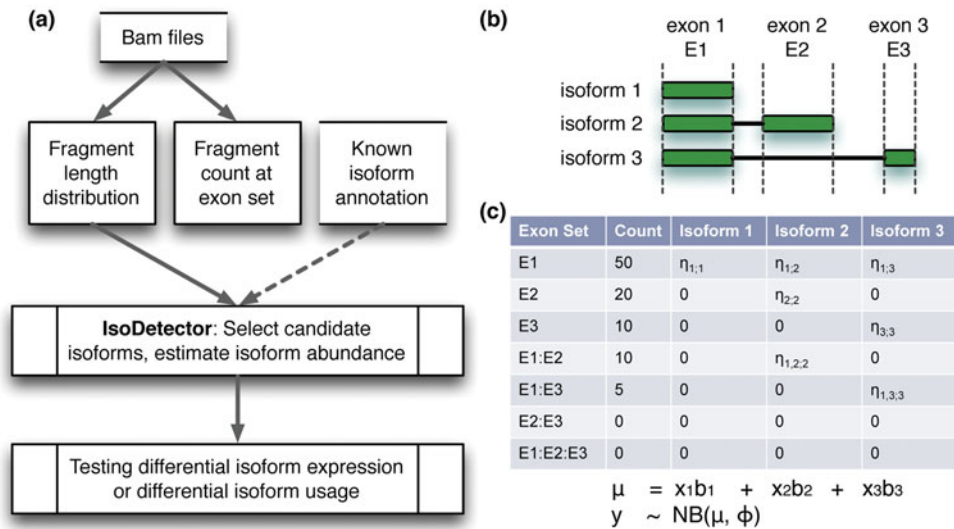


Figure 1. (a) IsoDOT work flow. The dashed line indicates that known isoform annotation (i.e., transcriptome annotation) is optional. (b) A gene with 3 exons and 3 possible isoforms. (c) A matrix of input data. Each row corresponding to an exon set. The column “Count” is the number of RNA-seq fragments at each exon set, and the columns “isoform k ” for $k = 1, 2, 3$ give the effective lengths of each exon set within each isoform, and specifically, $\eta_{A,k}$ is the effective length of exon set A for the k th isoform. $\text{NB}(\mu, \phi)$ indicates a negative binomial distribution with mean μ , and dispersion parameter ϕ .

2010; Li, Feng, and Jiang 2011b; Li et al. 2011a; Xia et al. 2011), IsoDetector employs penalized negative binomial regression with a log penalty. The negative binomial distribution is commonly used to model RNA-seq data, and previous studies have shown that it can account for variation in RNA-seq fragment counts across biological replicates (Langmead, Hansen, and Leek 2010). Many popular methods for differential expression testing, such as DESeq (Anders and Huber 2010) and edgeR (Robinson, McCarthy, and Smyth 2010), adopt such an assumption. More specifically, the negative binomial distribution assumption, denoted by $\text{NB}(\mu, \phi)$, implies that the RNA-seq fragment count across biological replicates follows a negative binomial distribution with mean value μ and variance $\mu + \mu^2\phi$, where ϕ is an over-dispersion parameter. Therefore, the variance of a negative binomial distribution can be arbitrarily large for a large value of ϕ . The log penalty, which can be interpreted as an iterative adaptive Lasso penalty (Tibshirani 1996; Zou 2006; Sun, Ibrahim, and Zou 2010), is flexible enough to handle a broad class of penalization problems (Chen, Sun, and Fine 2014). IsoDOT can test DIE/DIU against any categorical or continuous covariate at any sample size, with or without known isoform annotation. Our simulation and real data analysis demonstrate that IsoDOT performs very well with human or mouse RNA-seq data, but it can of course be applied to RNA-seq data from any species with a reference genome.

Two exons of a gene may overlap partially. In such situations, we split them into three exons: the part unique to the first or the second exon, and the part that belongs to both exons. Multiple genes may overlap on one or more exons, and we consider these genes as a transcript cluster, though any isoform of this transcript cluster can only be produced from one gene.

Isoform Estimation in a Single Sample

We study the isoforms of each transcript cluster separately, and the following discussion applies to a specific transcript clus-

ter. Denote the number of exons of a transcript cluster by k . Let A be an exon set, that is, a subset of the k exons. Let y_{iA} be the number of sequence fragments that overlap and only overlap with all the exons of A in the i th sample, where $1 \leq i \leq n$, and n is the sample size. A sequence fragment overlaps with an exon if the “sequenced portion” of this fragment overlaps with at least 1 bp of the exon. For example, if a fragment is sequenced by a paired-end read where the first end overlaps with exon 1 and 2 and the second end overlaps with exon 4, then this fragment is assigned to exon set $A = \{1, 2, 4\}$.

To illustrate the main feature of our method, we consider a gene (which is a transcript cluster itself) with 3 exons and 3 isoforms (Figure 1(b)). Denote its expression at sample i by $\mathbf{y}_i = (y_{i\{1\}}, y_{i\{2\}}, y_{i\{3\}}, y_{i\{1,2\}}, y_{i\{2,3\}}, y_{i\{1,3\}}, y_{i\{1,2,3\}})^T$. We assume y_{iA} follows a negative binomial distribution $\psi(\mu_{iA}, \phi)$ with unknown mean μ_{iA} and dispersion parameter ϕ . Let $\boldsymbol{\mu}_i$ be a column vector concatenating the μ_{iA} 's, then $\boldsymbol{\mu}_i = E(\mathbf{y}_i)$. We model $\boldsymbol{\mu}_i$ by

$$\boldsymbol{\mu}_i = \mathbf{X}_i \boldsymbol{\beta} = \sum_{u=1}^p \mathbf{x}_{iu} \beta_u, \quad \beta_u \geq 0, \quad (1)$$

where β_u is proportional to the transcript abundance of the u -th isoform, $\mathbf{X}_i = (\mathbf{x}_{i1}, \dots, \mathbf{x}_{ip})$, and \mathbf{x}_{iu} for $1 \leq u \leq p$ represents the effective lengths of all the exon sets for the u -th isoform in the i -th sample. Intuitively, *effective length* is the “usable length” for the data generation mechanism, that is, the number of positions where a randomly selected sequence fragment can be sampled. The effective length of an exon set varies across the underlying isoforms. For example, isoforms 1 and 3 of the gene shown in Figure 1(b)–(c) do not include exon 2, and thus, the effective length of exon set $\{2\}$ is 0 for isoforms 1 and 3. In contrast, the effective length of exon 2 is nonzero for isoform 2, since it includes exon 2. In addition, effective length is also a function of the sample-specific RNA-seq fragment length distribution. The

desired average length of RNA-seq fragments is often chosen during RNA-seq library preparation. However, the true fragment length distribution can be estimated from observed RNA-seq data. See Supplementary Materials Section A for details. In this example, the design matrix includes the effective lengths of all exon sets for isoforms 1, 2, and 3. Next, we recast the isoform estimation problem to a negative binomial regression problem with fragment counts \mathbf{y}_i as response and effective lengths \mathbf{X}_i as covariates:

$$\mathbf{y}_i \sim \psi(\boldsymbol{\mu}_i, \phi), \text{ and } \boldsymbol{\mu}_i = \mathbf{X}_i \boldsymbol{\beta}. \quad (2)$$

Equation (2) should be understood such that y_{iA} 's are independent of each other and y_{iA} follows a negative binomial distribution $\psi(\mu_{iA}, \phi)$. The independence assumption is reasonable because the fragment counts across exon sets should be independent given isoform configurations.

The regression problem presented in Equation (2) is challenging because there can be a large number of possible isoforms and their effective lengths (e.g., the columns of the design matrix \mathbf{X}_i) may be linearly dependent or significantly correlated. To address this difficulty, we first select a set of candidate isoforms, and then apply a penalized negative binomial regression to select the final set of isoforms from these candidate isoforms. The candidate isoforms can be selected using observed RNA-seq data (Supplementary Materials Section B) or a transcriptome annotation database (e.g., Ensembl; Flicek et al. 2011).

In our analysis, we skip the exon sets that have zero or negligible effective lengths across all the candidate isoforms because these exon sets are not informative for isoform expression estimation. For example, the exon set $\{2, 3\}$ or $\{1, 2, 3\}$ in the example shown in Figure 1(c) are not included in the analysis. The number of candidate isoforms, denoted by p , can be much larger than the number of (informative) exon sets, denoted by m , and there may be high correlations among the effective lengths of the candidate isoforms. Therefore, selecting the final set of isoforms from the candidate isoforms is a challenging variable selection problem. Lasso penalty has been applied in previous studies. However, the selection consistency of Lasso requires an *irrepresentability condition* on the design matrix (Zhao and Yu 2006; Zou 2006), which posits that there are weak correlations between the ‘‘important covariates,’’ which have nonzero effects and the ‘‘unimportant covariates,’’ which have zero effects. This irrepresentability condition is often not satisfied for the isoform selection problem due to high correlations among candidate isoforms. We employ a log penalty (Mazumder, Friedman, and Hastie 2011) for this challenging variable selection problem, which does not require the irrepresentability condition and can be interpreted as iterative adaptive Lasso (Sun, Ibrahim, and Zou 2010; Chen, Sun, and Fine 2014). The algorithm for fitting this penalized negative binomial regression is outlined in Supplementary Materials Section C.

Isoform Estimation in Multiple Samples

To estimate isoform expression in multiple samples, we have to account for read-depth difference across samples. Let t_i be a read-depth measurement for the i th sample. For example, t_i can be the total number of RNA-seq fragments in the i th sample, or the 75 percentile of the number of RNA-seq fragments per

gene in the i th sample (Bullard et al. 2010). We first consider a case without any covariate associated with isoform expression. To account for read-depth variation, we modify Equation (1) to

$$\boldsymbol{\mu}_i = t_i \mathbf{X}_i \boldsymbol{\gamma} = \sum_{u=1}^p t_i \mathbf{x}_{iu} \gamma_u, \quad (3)$$

where γ_u is proportional to relative expression of the u th isoform, after normalizing by t_i .

Let $\mathbf{y}^\top = (\mathbf{y}_1^\top, \dots, \mathbf{y}_n^\top)$, $\boldsymbol{\mu}^\top = (\boldsymbol{\mu}_1^\top, \dots, \boldsymbol{\mu}_n^\top)$, and $\mathbf{Z}^\top = (t_1 \mathbf{X}_1^\top, \dots, t_n \mathbf{X}_n^\top)$, where \mathbf{y} and $\boldsymbol{\mu}$ are vectors of length nm and \mathbf{Z} is a matrix of size $nm \times p$. Recall that p is the number of candidate isoforms, n is sample size, and m is the total number of exon sets. Then the isoform selection problem can be written as a negative binomial regression problem

$$\mathbf{y} \sim \psi(\boldsymbol{\mu}, \phi) \text{ and } \boldsymbol{\mu} = \mathbf{Z} \boldsymbol{\gamma}, \text{ where } \gamma_j \geq 0 \text{ for } 1 \leq j \leq p. \quad (4)$$

After imposing a penalty, we solve this regression problem using the method described in Supplementary Materials Section C.

Next we consider isoform estimation given a continuous covariate $\mathbf{g} = (g_1, \dots, g_n)^\top$. We assume the expression of an isoform u for sample i , denoted by γ_{iu} , is a linear function of covariate g_i : $\gamma_{iu} = a_u + \tilde{b}_u g_i$. This linear model is an appropriate choice when g_i represents SNP genotype (Sun 2012), which is the focus of our empirical data analysis. In this linear model setup, a complex set of constraints is needed for a_u and \tilde{b}_u so that $\gamma_{iu} \geq 0$ for any value of g_i . Therefore, we reformulate the problem as follows. Without loss of generality, we scale the value of g_i to be within the range of $[0, 1]$ with the minimum and maximum values being exactly 0 and 1, respectively. For example, if g_i corresponds to a SNP with additive effect, we can set $g_i = 0, 0.5$, or 1 for genotype AA, AB, or BB. Let $b_u = \tilde{b}_u + a_u$, then $\gamma_{iu} = a_u + (b_u - a_u)g_i = a_u(1 - g_i) + b_u g_i$. Under this model, we have

$$\gamma_{iu} \geq 0 \text{ for any } g_i \in [0, 1] \Leftrightarrow a_u \geq 0 \text{ and } b_u \geq 0.$$

Let $\mathbf{a} = (a_1, \dots, a_p)^\top$ and $\mathbf{b} = (b_1, \dots, b_p)^\top$, we have

$$\boldsymbol{\mu}_i = t_i \mathbf{X}_i [\mathbf{a}(1 - g_i) + \mathbf{b}g_i]. \quad (5)$$

By concatenating \mathbf{a} and \mathbf{b} into a vector: $\boldsymbol{\theta} = (a_1, \dots, a_p, b_1, \dots, b_p)^\top$, we can rewrite Equation (5) as $\boldsymbol{\mu}_i = \mathbf{W}_i \boldsymbol{\theta}$, where $\mathbf{W}_i = [t_i \mathbf{X}_i (1 - g_i), t_i \mathbf{X}_i g_i]$ is an $m \times 2p$ matrix. Let $\mathbf{W}^\top = (\mathbf{W}_1^\top, \dots, \mathbf{W}_n^\top)$, then the isoform expression estimation problem reduces to a negative binomial regression problem with nonnegative coefficients

$$\mathbf{y} \sim \psi(\boldsymbol{\mu}, \phi) \text{ and } \boldsymbol{\mu} = \mathbf{W} \boldsymbol{\theta}, \text{ where } \theta_j \geq 0 \text{ for } 1 \leq j \leq 2p. \quad (6)$$

After imposing a penalty, we solve the resulting penalized regression problem by the coordinate ascend method described in Supplementary Materials Section C.

Finally, we consider the general situation with q covariates, denoted by $\mathbf{g}_1, \dots, \mathbf{g}_q$, where $\mathbf{g}_v = (g_{1v}, \dots, g_{nv})^\top$ for $v = 1, \dots, q$. Without loss of generality, we assume $0 \leq g_{iv} \leq 1$ for $1 \leq i \leq n$ and $1 \leq v \leq q$. Then we model γ_{iu} by $\gamma_{iu} = qa_u + \sum_{v=1}^q (b_{vu} - a_u)g_{iv} = a_u \sum_{v=1}^q (1 - g_{iv}) + \sum_{v=1}^q b_{vu} g_{iv}$. This is simply a multiple linear regression model where each covariate g_v has its own effect. Let

$\mathbf{a} = (a_1, \dots, a_p)^\top$ and $\mathbf{b}_v = (b_{1v}, \dots, b_{pv})^\top$, then,

$$\boldsymbol{\mu}_i = t_i \mathbf{X}_i \left[\mathbf{a} \sum_{v=1}^q (1 - g_{iv}) + \mathbf{b}_1 g_{i1} + \dots + \mathbf{b}_q g_{iq} \right]. \quad (7)$$

By concatenating \mathbf{a} , \mathbf{b}_1 , $\mathbf{b}_2, \dots, \mathbf{b}_q$ into a vector: $\boldsymbol{\theta} = (\mathbf{a}^\top, \mathbf{b}_1^\top, \dots, \mathbf{b}_q^\top)^\top$, we rewrite the above equation as $\boldsymbol{\mu}_i = \mathbf{W}_i \boldsymbol{\theta}$, where $\mathbf{W}_i = [t_i \mathbf{X}_i \sum_{v=1}^q (1 - g_{iv}), t_i \mathbf{X}_i g_{i1}, \dots, t_i \mathbf{X}_i g_{iq}]$ is an $m \times (q + 1)p$ matrix. Let $\mathbf{W}^\top = (\mathbf{W}_1^\top, \dots, \mathbf{W}_n^\top)$, then we form a negative binomial regression problem

$$\mathbf{y} \sim \psi(\boldsymbol{\mu}, \phi) \text{ and } \boldsymbol{\mu} = \mathbf{W}\boldsymbol{\theta}, \text{ for } \theta_j \geq 0, 1 \leq j \leq (q + 1)p. \quad (8)$$

After imposing a penalty on each θ_j , we can solve the resulting penalized regression problem by the coordinate ascend method described in Supplementary Materials Section C.

When studying multiple samples, it is possible that an exon set is only expressed in a subset of samples. Therefore, when examining the fragment counts of this exon set across all samples, there are more 0's than expected by a negative binomial distribution. In such case, we introduce a zero-inflated component and employ the zero-inflated negative binomial distribution (Rashid et al. 2011) to model RNA-seq fragment count data.

Testing for Differential Isoform Expression (DIE)

We have described the statistical model to estimate RNA isoform expression in multiple samples given one or more covariate. Building on this model, we assess DIE with respect to a set of covariates denoted by V using a likelihood ratio (LR) test. Specifically, the null hypothesis (H_0) is that $b_{uv} = a_u$ for $u = 1, \dots, p$ and $v \in V$ and the alternative hypothesis (H_1) is that $b_{uv} \neq a_u$ for at least one pair of (u, v) , where $u = 1, \dots, p$ and $v \in V$. It is helpful to understand this test by considering two special cases. First, we assume there is only one numerical covariate. Under H_0 , we solve the isoform estimation problem by a penalized negative binomial regression with expected value $\boldsymbol{\mu}_i = t_i \mathbf{X}_i \mathbf{a}$. Under alternative, $\boldsymbol{\mu}_i = t_i \mathbf{X}_i [\mathbf{a}(1 - g_i) + \mathbf{b}g_i]$. Therefore, the number of parameters is p under H_0 and $2p$ under H_1 . The asymptotic chi-square distribution with degree of freedom p does not apply because the models are estimated, under both H_0 and H_1 , by penalized regression. In the second special case, we assume the only variable of interest is a categorical variable with d categories. This categorical variable can be coded as $d - 1$ binary variables, denoted by $g_{i1}, \dots, g_{i,d-1}$. The expected values of fragment counts across exon sets under H_0 and H_1 are $\boldsymbol{\mu}_i = (d - 1)t_i \mathbf{X}_i \mathbf{a}$ and $\boldsymbol{\mu}_i = t_i \mathbf{X}_i [\mathbf{a} \sum_{v=1}^{d-1} (1 - g_{iv}) + \mathbf{b}_1 g_{i1} + \dots + \mathbf{b}_{d-1} g_{i,d-1}]$, respectively. Therefore, the number of parameters under H_0 and H_1 are p and dp , respectively. Again, the asymptotic chi-square distribution with degree of freedom $(d - 1)p$ does not apply because the models are estimated, under both H_0 and H_1 , by penalized regression.

We use LR statistic as our test statistic:

$$\mathcal{LR} = 2(\ell_1 - \ell_0), \quad (9)$$

where ℓ_0 and ℓ_1 are the log-likelihoods under null and alternative, respectively. Because of penalized estimation, the null distribution of this test statistic no longer follows the standard asymptotic distribution. We obtain the null distribution by parametric bootstrap or permutation. The parametric bootstrap ap-

proach proceeds as follows. (1) Fit the penalized negative binomial regression under null. (2) Sample fragment counts based on the fitted null model. (3) Using the sampled counts to refit models under null and alternative and calculate the LR statistic. (4) Repeat Steps (1)–(3) a large number of times as needed (Jiang and Salzman 2012). (5) Calculate the p -value as the proportion of the bootstrapped LR statistics that are larger than \mathcal{LR} . This parametric bootstrap procedure yields valid p -values regardless of the sample size n . However, since the sampling population is all the RNA-seq fragments from the studied samples, small p -values only imply a significant difference of the studied samples and should not be generalized to other individuals.

If sample size is sufficiently large (e.g., ≥ 5 cases vs. ≥ 5 controls), we can obtain valid statistical inference for the corresponding population (instead of studied samples) by calculating permutation p -values. Specifically, the null model is fitted without the covariate of interest, and thus, its log-likelihood ℓ_0 remains unchanged across permutations. In each permutation, we permute the covariate of interest and refit the alternative model, and then calculate the LR test statistic. We repeat this process a large number of times to obtain a null distribution of the LR statistic.

Testing for Differential Isoform Usage (DIU)

All the previous discussion, including RNA isoform expression estimation and DIE testing, focus on absolute expression of RNA isoforms, which is RNA isoform expression after correcting for read-depth variation across samples. Another measure of RNA isoform expression, which may be more interesting in many situations, is the relative expression with respect to the total expression of the corresponding gene or transcript cluster. This is because a gene may have higher or lower expression overall, and it may also switch the usage of some RNA isoforms. For example, a gene may predominately use one isoform in one tissue and switch to another isoform in another tissue. Such relative expression of an RNA isoform is referred to the isoform usage. For a transcript cluster of interest, we denote the total number of RNA-seq fragments in the i -th sample by r_i . Recall that in the previous discussions of DIE, t_i denotes a read-depth measurement for the i th sample. To assess DIU, we apply the same procedure as assessing DIE, except that we replace t_i by r_i .

RESULTS AND DISCUSSIONS

Simulation for Case-Control comparison

We simulated ~ 1 million 76 + 76 bps paired-end RNA-seq reads for a single case and control sample, respectively, using Flux Simulator (<http://flux.sammeth.net/simulator.html>) and the Ensembl transcriptome annotations (version 67, <http://useast.ensembl.org/info/data/ftp/index.html>) for chromosome 1 (chr1) and chromosome 18 (chr18) for the mouse genome. We simulated the data such that all genes from chr1 were equivalently expressed and all genes from chr18 were differentially expressed, either in terms of total expression or isoform usage (Supplementary Figure 2). These simulated RNA-seq reads were mapped to the reference genome using Tophat (Trapnell, Pachter, and Salzberg 2009). Next, we counted the

number of RNA-seq fragments per exon set. It was important to consider all exon sets rather than just exon or exon junctions because many RNA-seq fragments overlap with more than two exons. For example, in this simulated dataset, $\sim 27\%$ of the paired-end reads overlapped 3 or more exons (Supplementary Figure 3). In addition, we confirmed that the number of sequence fragments per exon set was proportional to the effective length of the exon set (Supplementary Figure 4).

The dimension of this problem (i.e., the number of exon sets m versus the number of candidate isoforms p) was illustrated in Supplementary Figures 5–6. Given transcriptome annotation, $p < m$ for the vast majority of transcript clusters, and without transcriptome annotation, we restricted the number of candidate isoforms so that approximately $p < 10m$. In either case, there were strong correlations among the candidate isoforms (Supplementary Figures 7–8), therefore, necessitating the use of penalized regression in IsoDetector. After applying penalized regression, most transcription clusters included 10 or fewer isoforms, with or without transcriptome annotation (Supplementary Figures 9–10).

We compared isoform abundance estimates from IsoDetector and Cufflinks (v2.0.0) (Trapnell et al. 2010, 2013) using the case sample. The conclusions from the control sample were similar (results not shown). We focused on 1062 transcript clusters (corresponding to 5185 transcripts) that had at least 2 exons with ≥ 5 sequence fragments overlapping each exon. Most of these transcript clusters included 1–2 genes and 1–42 transcripts, and most of the 5185 transcripts harbored 6–500 RNA-seq fragments (Figure 2(a)). If transcriptome annotation was unavailable, the isoforms selected by IsoDetector were more similar to the true ones than Cufflinks (Figure 2(b)). The two methods had similar accuracy in terms of transcript abundance estimation, either with (Figure 2(c)–(d)) or without (Figure 2(e)–(f)) transcriptome annotation.

We next compared the power of IsoDOT to that of Cuffdiff (v2.0.0) (Trapnell et al. 2013) in terms of testing for differential expression or DIU. Cuffdiff results were obtained from three files: `gene_exp.diff` (differential expression), `splicing.diff` (differential usage of the isoforms sharing a transcription start site (TSS)), and `promoters.diff` (differential usage of TSSs). The majority of the genes in file `gene_exp.diff` have status “OK” and they were used in the following comparison. However, no gene has status “OK” in file `splicing.diff` or `promoters.diff`. In these two files, all the genes with meaningful p -values have status “NOTEST,” indicating that Cuffdiff recommends users not to trust these testing results. The reason is as follows. For a two group comparison with one case and one control, Cuffdiff combines the case and the control to estimate variance of biological replicates, which leads to a very conservative p -value since this approach implicitly assumes the case and the control have the same expected value. Nonetheless, these genes with status “NOTEST” were used in the following comparison because they are the only genes that we can use. A gene could have multiple p -values in `splicing.diff`. In favor of Cuffdiff for power comparison, we used the smallest p -value for each gene. Based on our simulation setup, power was defined as the proportion of the genes from chr18 that had significant DIE or DIU. IsoDOT had substantially higher power than Cuffdiff (Figure 3), attributed to the correct

Type I error control of IsoDOT (Supplementary Figure 11). The poor performance of Cuffdiff is not simply a calibration issue because IsoDOT still performs better than Cuffdiff when we compare two methods using ROC curves (Supplementary Figure 12). Instead, the poor performance of Cuffdiff is because it tries to estimate variance across biological replicates when there is no biological replicate at all. IsoDOT can perform a valid test because it does not try to estimate variance of biological replicates. Instead, it tries to estimate the variance due to resampling of RNA-seq reads. To be fair, Cuffdiff does recommend such test with one case versus one control and we include Cuffdiff in our comparison because it is the only method that allows for such testing. In conclusion, this simulation illustrated that IsoDOT worked well in this challenging situation of comparing one case versus one control.

Simulation for Isoform Usage eQTL

To illustrate the performance of IsoDOT in testing DIU with respect to a continuous covariate, we applied IsoDOT to map isoform usage eQTL (gene expression quantitative trait locus) using simulated RNA-seq data and real genotype data. We downloaded genotype data from 60 European HapMap samples (Thorrisson et al. 2005) and selected 949,537 SNPs with minor allele frequency (MAF) > 0.05 for the following analysis. We defined transcription clusters based on Ensembl annotation (version 66, <http://useast.ensembl.org/info/data/fjp/index.html>), and selected 200 transcript clusters that satisfied the following two conditions for our simulation studies. (1) Each transcript cluster has > 1 annotated RNA-isoforms, and (2) Each transcript cluster has ≥ 1 SNP in the gene body (any intronic or exonic regions) or within 1000bp of the gene body. For each selected transcript cluster, we simulated RNA-seq fragment counts across exon sets under null (H_0) and alternative (H_1), respectively. Specifically, for each gene, we randomly selected 50% of the isoforms to have zero expression and set the expression of the other 50% of isoforms by drawing γ_u (Equation (4)) from a uniform distribution $U[0.5, 1]$.

Next, we used these simulated data to assess DIU of each transcript cluster with respect to each of the nearby SNPs (within 1000bp of the gene body) and kept the most significant eQTL per transcript cluster. For each transcript cluster, up to 1000 permutations were carried out to correct for multiple testing across the multiple nearby SNPs. Under H_0 , the permutation p -values followed a uniform distribution; and under H_1 , the permutation p -values were obviously enriched with small values (Figure 4(a)). In this simulation, we had $\sim 80\%/40\%$ power to detect local isoform usage eQTL for permutation p -value cutoffs 0.05 and 0.005, respectively (Figure 4(b)). Therefore, this simulation demonstrates that IsoDOT correctly controls Type I error and has power to detect DIU with respect to a continuous covariate.

Haloperidol Treatment Effect on Mouse Transcriptome

Haloperidol is a drug used to treat schizophrenia, acute psychotic states, and delirium. A major adverse side effect of chronic haloperidol treatment is tardive dyskinesia (TD). TD can be modeled in mice by examining haloperidol-induced vacuous chewing movements (VCMs) (Crowley et al. 2012).

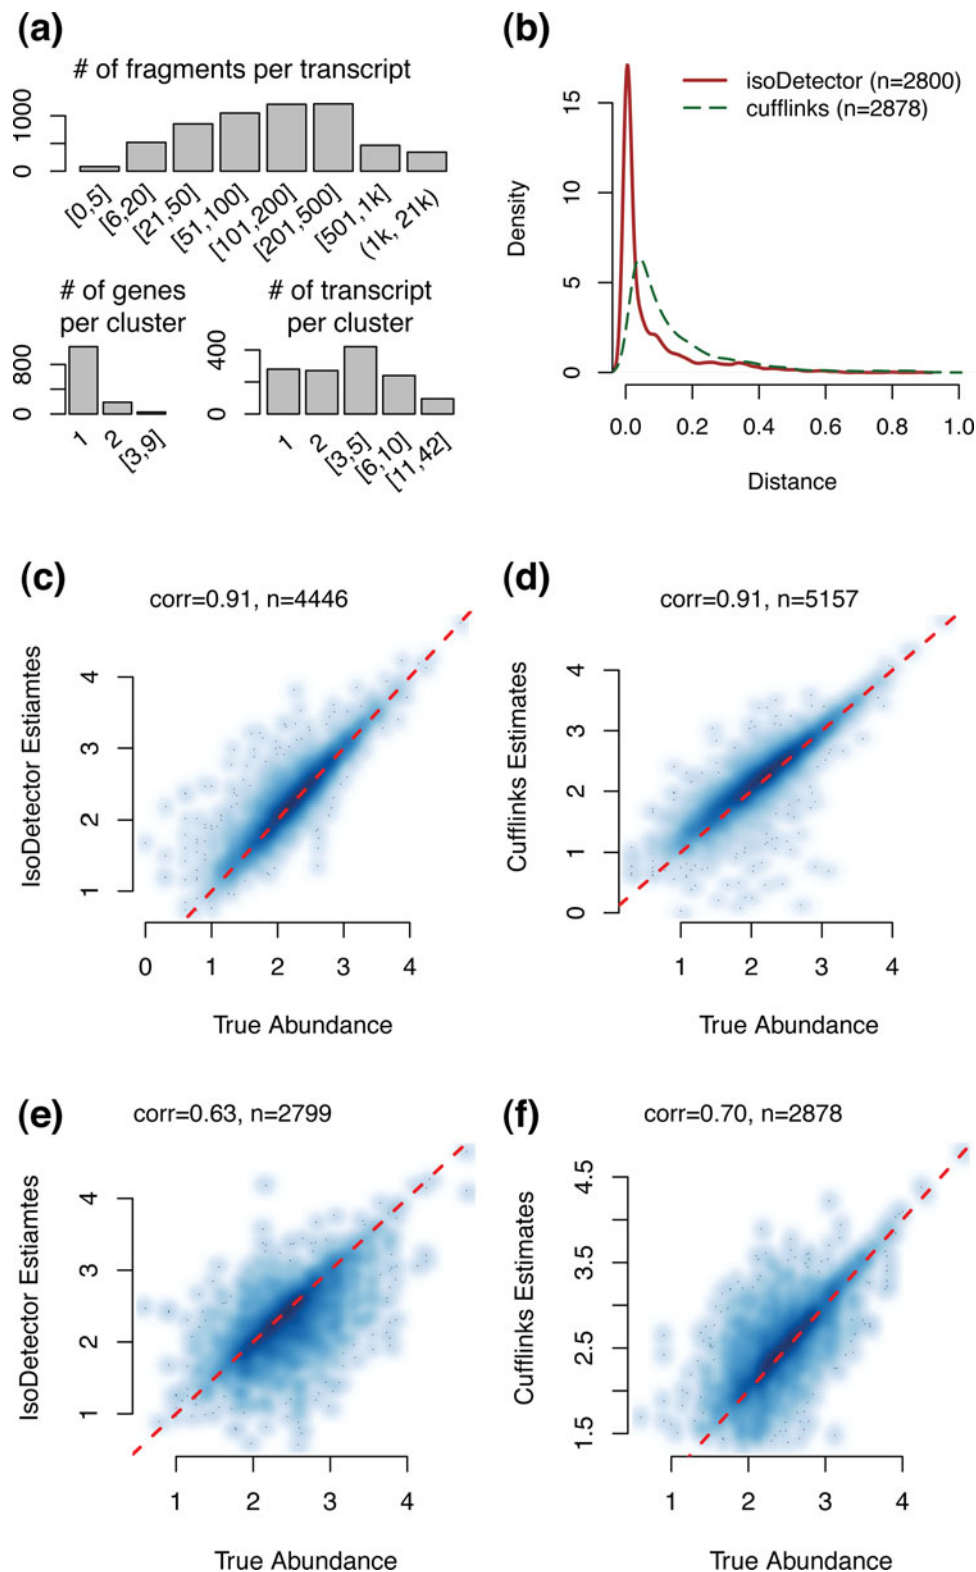


Figure 2. (a) A summary of the RNA-seq data and annotation of the simulated case sample. (b) Density curves of the distance between each de novo transcript and its closest transcript from transcriptome annotation. The distance is defined as the ratio of the number of unmatched base pairs over the total number of base pairs covered by either isoform. A base pair is “matched” if it corresponds to an exon or intron location for both isoforms. (c)–(d) Comparison of true transcript abundance versus estimates from IsoDetector or Cufflinks when we use known isoform annotation. Both X-axis and Y-axis are in log₁₀ scale. “n” is the number of transcripts with status “OK” for either IsoDetector or Cufflinks. (e)–(f) Comparison of true transcript abundance versus estimates from IsoDetector or Cufflinks when we do not use any isoform annotation.

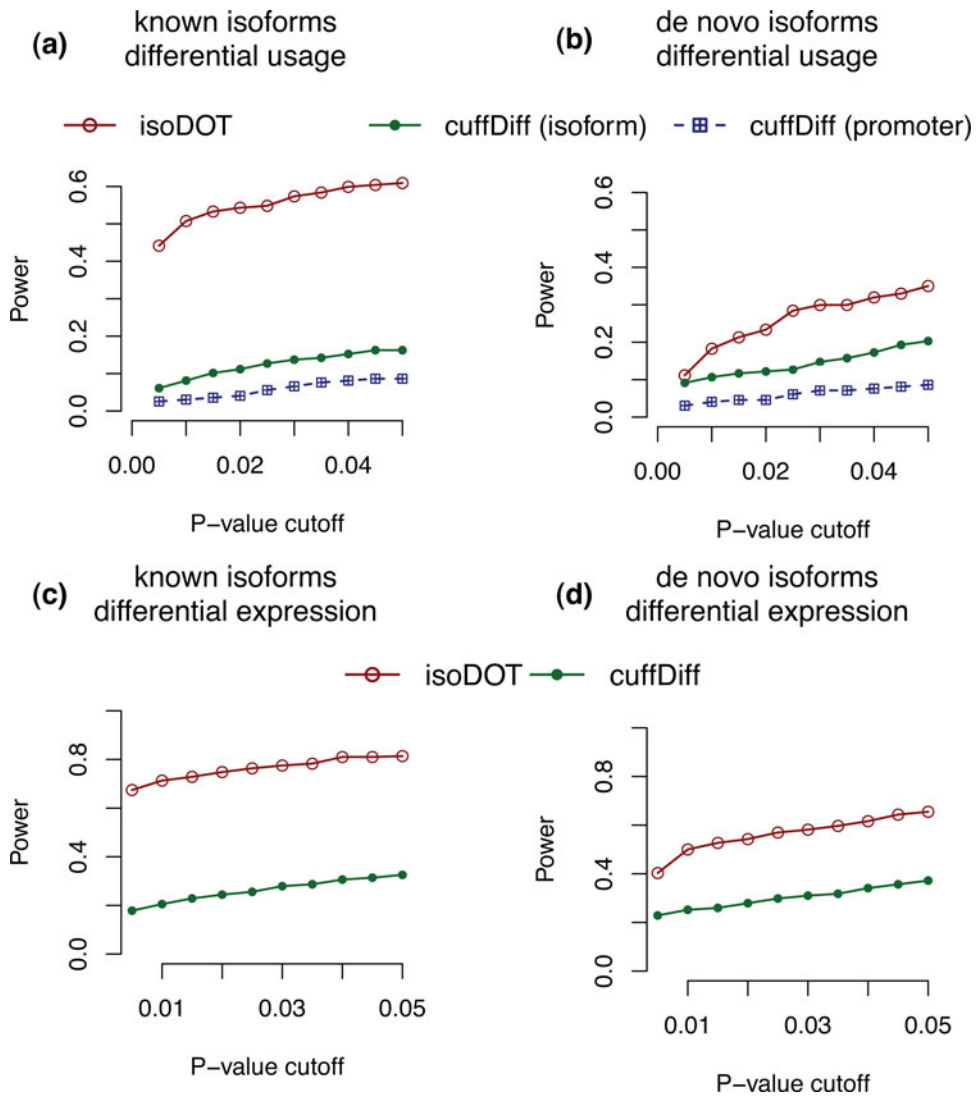


Figure 3. Compare the power of IsoDOT and Cuffdiff for detecting genes with differential isoform usage (a)–(b) or differential isoform expression (c)–(d), while transcriptome annotation is known (a)–(c) or not (b)–(d).

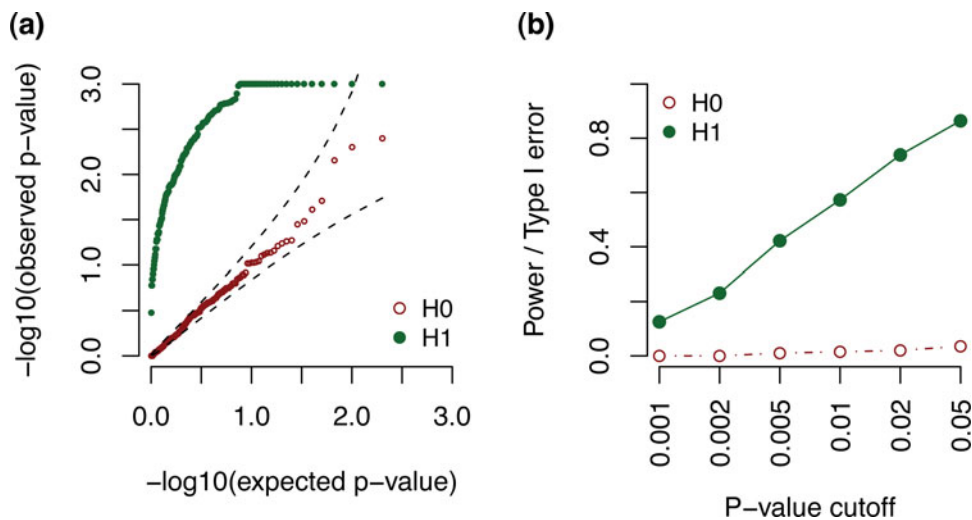


Figure 4. Simulation results for testing isoform-usage eQTL on 200 transcription clusters. (a) The qq-plot for p -values distributions against expected uniform distribution. (b) Power (H1) or Type I error (H0) for different p -value cutoffs.

We applied IsoDOT to analyze RNA-seq data for mice treated with haloperidol versus placebo with particular interest in identifying genes responding to haloperidol treatment and/or responsible for VCM. RNA-seq data were collected from whole brains from four mice: two C57BL/6J mice treated with haloperidol or placebo and two 129S1Sv/ImJ×PWK/PhJ F1 mice treated with haloperidol or placebo. Each RNA-seq fragment was sequenced on both ends by 93 or 100bp, and 20–27 million RNA-seq reads were collected for each mouse (Supplementary Table 1). See Supplementary Materials Section D for additional details of the experiment.

We first studied DIE/usage between two C57BL/6J mice. RNA-seq reads were mapped to the mm9 reference genome using TopHat (Trapnell, Pachter, and Salzberg 2009). At FDR of 5%, IsoDOT identified 86 or 88 genes with DIU, with or without transcriptome annotation, respectively. For the test of DIE, also at FDR of 5%, IsoDOT identified 332 or 206 genes with or without transcriptome annotation, respectively. We sought to gain some insight of these four gene lists by applying functional category enrichment analysis (Sherman et al. 2009) on the top 100 genes in each list. Only those DIU genes identified with transcriptome annotation were significantly associated with biologically relevant categories such as neuron projection (Supplementary Figures 13–16). This implied that DIU, rather than DIE, might be more relevant to haloperidol treatment in C57BL/6J mice. The gene lists identified without transcriptional annotation did not show functional enrichment, which implied larger sample size or higher read-depth were needed to detect DIE or DIU without transcriptome annotation. According to Cuffdiff manual, it is not recommended to run Cuffdiff with sample size one versus one. However, we still evaluated the performance of Cuffdiff in this dataset for comparison purpose. Using the results reported by Cuffdiff, no gene has q -value smaller than 0.05 (with or without annotation, alternative promoter or alternative splicing), and functional category enrichment analysis (Sherman et al. 2009) on the top 100 genes reported by Cuffdiff identified no significantly enriched functional category (Supplementary Figure 17).

Several DIU genes identified by IsoDOT can be potential targets for follow up studies (Supplementary Table 2). For example, *Utrn* (utrophin, Figure 5) and *Dmd* (dystrophin) are both involved in neuron projection and could be candidates underlying the VCM side effect. *Grin2b* (glutamate receptor, ionotropic, NMDA2B, Supplementary Figure 18) or its human ortholog is involved in Alzheimer's disease, Huntington's disease, and amyotrophic lateral sclerosis (ALS), which are potentially relevant to the VCM side effect of haloperidol treatment. In addition, our previous studies had prioritized several other glutamate receptors such as *Grin1* and *Grin2a* as candidates that response to haloperidol treatment using independent data and methods (Crowley et al. 2012).

Similar studies were conducted for the two F1 mice of 129S1Sv/ImJ×PWK/PhJ. To better map RNA-seq reads, we first built two pseudogenomes for 129S1Sv/ImJ and PWK/PhJ by incorporating Sanger SNPs and indels (Keane et al. 2011) into the reference genome and mapped RNA-seq reads to the two pseudogenomes separately. Next, alignments were remapped back to the reference coordinate system and the observed ge-

netic variants were annotated for each RNA-seq fragment. IsoDOT identified much less DIU or DIE genes in these two F1 mice than in the two C57BL/6J mice. At FDR 0.2, no DIU genes were identified and 85 DIE gene were identified. Six of these 85 genes were involved in actin-binding, though this functional category was not significantly overrepresented.

The greater level of DIE and DIU in C57BL/6J than 129S1Sv/ImJ×PWK/PhJ following chronic haloperidol treatment was consistent with behavioral phenotype data which showed that C57BL/6J mice had greater susceptibility to haloperidol-induced VCM (Supplementary Figure 19) (Crowley et al. 2012). Therefore, some of the C57BL/6J transcriptional changes detailed in this study might contribute to the development of haloperidol-induced VCM.

Allele-Specific Differential Isoform Usage

About 37.2% of the RNA-seq fragments from the two 129S1Sv/ImJ×PWK/PhJ F1 mice were only mapped to the paternal or maternal allele or were mapped to one allele with fewer mismatches, and hence, they were allele-specific RNA-seq fragments. There was no genome-wide bias, that is, ~50% of the allele-specific RNA-seq fragments were mapped to each parental strain (Supplementary Table 1). Using these allele-specific RNA-seq fragments, we applied IsoDOT to assess DIU between maternal and paternal alleles. At a liberal FDR cutoff 0.2, no DIU gene was identified, and 19 or 30 DIE genes were identified from the haloperidol/placebo treated F1 mice, respectively. The genes with significant DIU in the haloperidol treated mice, but not the placebo treated mice might indicate genetic×treatment interactions. Supplementary Table 3 listed 23 such genes with DIU p -values < 0.01 in the haloperidol treated mice, and DIU p -values > 0.1 in the placebo treated mice. Among them, *Synpo* and *Snap25* are associated with neuron functions. *Snap25* is known to be associated with schizophrenia and/or haloperidol treatment at DNA (Müller et al. 2005), RNA (Sommer et al. 2010), and protein levels (Thompson, Sower, and Perrone-Bizzozero 1998; Honer et al. 2002). To the best of our knowledge, this is the first report that the DIU of *Snap25* is associated with genetic×haloperidol treatment interaction.

Software and Data Availability

An R package of IsoDOT is available at <http://www.bios.unc.edu/~weisun/software/isoform.htm>. Testing DIE/usage is computationally intensive. Using IsoDOT with up to 1000 parametric bootstraps, it will take on average 1–3 minutes to test DIU for a gene on a single processor. Parallel computation is needed and straightforward for genome-wide study. We are also actively working on implementing our method using graphics processing unit (GPU) using the massively parallel algorithm described in Supplementary Materials Section E. Simulated RNA-seq data can be downloaded from http://www.bios.unc.edu/~weisun/software/isoform_files/. The RNA-seq data of mouse haloperidol treatment experiment are available from NCBI GEO.

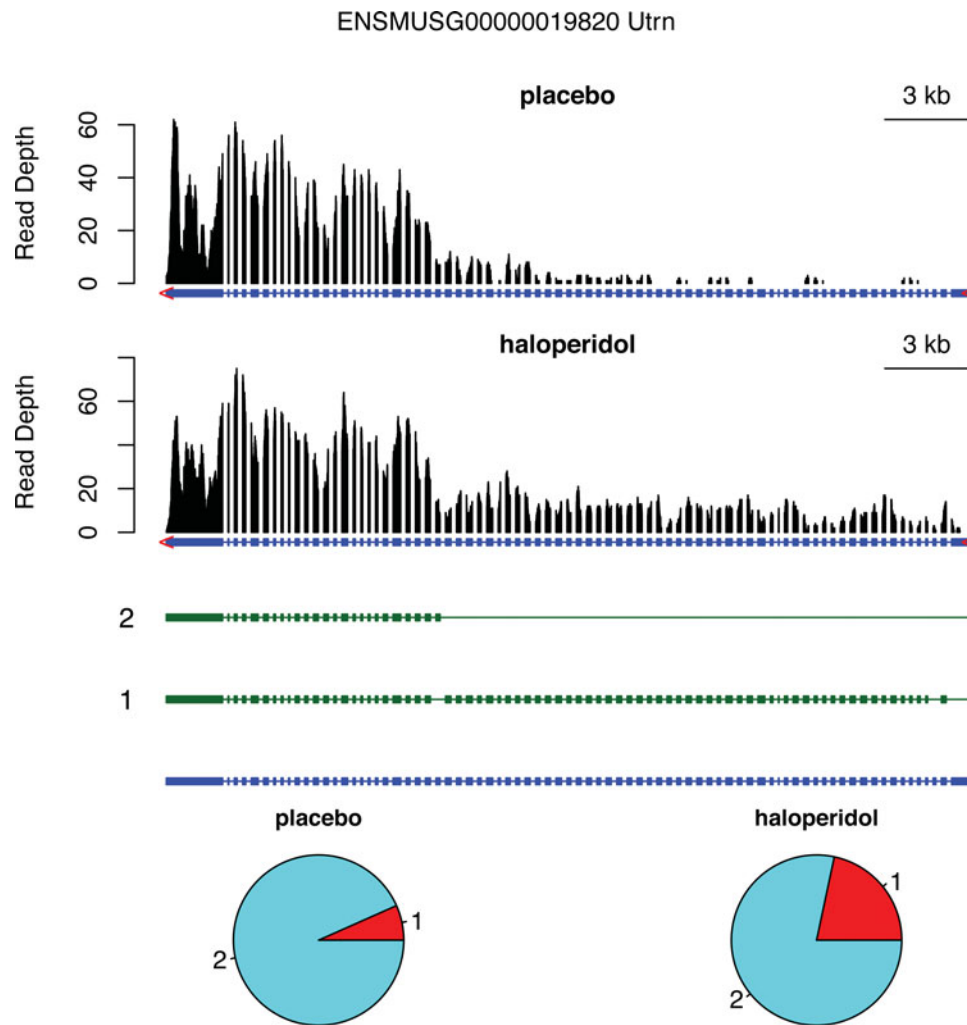


Figure 5. Differential isoform usage of gene Utrn between a mouse with haloperidol treatment and a mouse with placebo. The top panel shows the read-depth of RNA-seq reads in two conditions. Two annotated isoforms and all the exons are illustrated in middle panel. The expression of each isoform under two conditions are shown in the bottom panel.

Discussion

We have developed a new statistical method named IsoDOT to assess DIE or usage from RNA-seq data with respect to categorical or continuous covariates. The resampling based approach is the basis of our hypothesis testing method. Two components of our method, the negative binomial distribution assumption and the log penalty for penalized estimation are important for the success of this resampling based approach. First, the negative binomial distribution is a well-accepted choice for modeling RNA-seq fragment data across biological replicates. For the completeness of our article, we also demonstrate that the negative binomial distribution provides a good fit of RNA-seq fragment counts, while the Poisson distribution assumption leads to a severe underestimate of variance (Supplementary Figure 20). Replacing the log penalty with the Lasso penalty in IsoDOT leads to inaccurate Type I error control and/or reduced power (Supplementary Figure 21). The limitation of Lasso is especially apparent when we do not use isoform annotation, where the number of candidate isoforms is much larger than sample size. This is consistent with previous findings that Lasso tends to select more false positives and has larger

bias in effect estimates than the log penalty (Chen, Sun, and Fine 2014).

Some biases should be accounted for to obtain better estimates of RNA isoform abundance. For example, the RNA-seq reads may not be uniformly distributed along the transcript, and DNA sequence features such as GC content may affect the abundance of RNA-seq reads. Such biases, if they exist, affect both the likelihoods under the null and alternative hypotheses and do not alter our Type I error rate. Furthermore, it may have limited impact on power because we calculate p -values through resampling approaches. For example, if the last exon of a gene tends to have a larger number of RNA-seq reads, our method may overestimate the expression of the isoform harboring the last exon. However, such over-estimation occurs under both the null and alternative hypotheses and, therefore, does not lead to inflated or deflated Type I error. Systematically accounting for such biases are among our future plans to improve IsoDOT's performance.

A major strength of IsoDOT is that it allows the assessment of DIU between one case and one control sample. This is especially important for paired samples, for example, maternal and paternal alleles of one individual. When there are multiple

paired samples, we can combine the p -values of multiple pairs using meta-analysis via Fisher's method or Stouffer's Z-score (Hunter, Schmidt, and Jackson 1982). In the preliminary study reported in this article, we have identified several interesting genes (*Utrn*, *Dmd*, *Grin2b*, *Snap25*) whose isoform usage may respond to haloperidol treatment and are perhaps related to its side effects. In the near future, we plan to extend this study to include a much larger number of mice with diverse genetic backgrounds including mice from the Collaborative Cross (Churchill et al. 2004; Consortium 2012).

Recently developed sequencing techniques can now deliver longer reads, including 2×250 bp reads from Illumina's MiSeq or 400 bp reads from Ion Torrent. Our methods can handle these longer sequencing reads without any difficulty. In fact, when the RNA-seq reads are long enough, transcriptome reconstruction becomes much easier. However, until all the isoforms of a transcript cluster can be unambiguously reconstructed and all the RNA-seq fragments can be assigned to an isoform with (almost) 100% certainty, testing DIE/usage remains a challenging problem where methods and software such as IsoDOT are needed. Another recently developed RNA-seq technique is to deliver "stranded" sequences so that RNA-seq from sense and anti-sense strands are separated. For the analysis of such stranded RNA-seq data, the only step in the IsoDOT pipeline that needs to be modified is to count the RNA-seq fragments for sense and anti-sense strands separately.

SUPPLEMENTARY MATERIALS

Section A: Calculation of effective length.

Section B: Selection of candidate isoforms.

Section C: Model fitting of the penalized negative binomial regression.

Section D: Mouse haloperidol treatment experiment.

Section E: Massively parallel computing for IsoDetector.

Section F: Additional figures.

[Received May 2013. Revised February 2015.]

REFERENCES

- Alamancos, G. P., Agirre, E., and Eyraes, E. (2014), "Methods To Study Splicing From High-Throughput RNA Sequencing Data," in *Spliceosomal Pre-mRNA Splicing*, New York: Springer. pp. 357–397. [975]
- Anders, S., and Huber, W. (2010), "Differential Expression Analysis for Sequence Count Data," *Genome Biology*, 11, R106. [977]
- Anders, S., Reyes, A., and Huber, W. (2012), "Detecting Differential Usage of Exons From RNA-Seq Data," *Genome Research*, 22, 2008–2017. [976]
- Barbosa-Morais, N. L., Irimia, M., Pan, Q., Xiong, H. Y., Gueroussov, S., Lee, L. J., Slobodenic, V., Kutter, C., Watt, S., Çolak, R., Kim, T., Misquitta-Ali, C. M., Wilson, M. D., Kim, P. M., Odom, D. T., Frey, B. J., and Blencowe, B. J. (2012), "The Evolutionary Landscape of Alternative Splicing in Vertebrate Species," *Science*, 338, 1587–1593. [975]
- Bohner, R., and Rätsch, G. (2010), "Rquant: Web: A Tool for RNA-Seq-Based Transcript Quantitation," *Nucleic Acids Research*, 38, W348–W351. [975,977]
- 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. [978]
- Chen, L. (2012), "Statistical and Computational Methods for High-Throughput Sequencing Data Analysis of Alternative Splicing," *Statistics in Biosciences*, 5, 138–155. [975]
- Chen, T. H., Sun, W., and Fine, J. (2014), *Designing Penalty Functions in High Dimensional Problems: The Role of Tuning Parameters*, Chapel Hill, NC: UNC Press. [977,978,984]
- Churchill, G. A., Airey, D. C., Allayee, H., Angel, J. M., Attie, A. D., Beatty, J., Beavis, W. D., Belknap, J. K., Bennett, B., Berrettini, W. et al. (2004), "The Collaborative Cross, a Community Resource for the Genetic Analysis of Complex Traits," *Nature Genetics*, 36, 1133–1137. [985]
- Consortium, C. C. (2012), "The Genome Architecture of the Collaborative Cross Mouse Genetic Reference Population," *Genetics*, 190, 389–401. [985]
- Crowley, J. J., Adkins, D. E., Pratt, A. L., Quackenbush, C. R., van den Oord, E. J., Moy, S. S., Wilhelmsen, K. C., Cooper, T. B., Bogue, M. A., McLeod, H. L. and Sullivan, P. F. (2012), "Antipsychotic-Induced Vacuolar Chewing Movements and Extrapyramidal Side Effects are Highly Heritable in Mice," *The Pharmacogenomics Journal*, 12, 147. [980,983]
- Denoeud, F., Aury, J. M., Da Silva, C., Noel, B., Rogier, O., Delledonne, M., Morgante, M., Valle, G., Wincker, P., Scarpelli, C., Jaillon, O., and Artiguenave, F. (2008), "Annotating Genomes With Massive-Scale RNA Sequencing," *Genome Biology*, 9, R175. [975]
- Djebali, S., Davis, C. A., Merkel, A., Dobin, A., Lassmann, T., Mortazavi, A., Tanzer, A., Lagarde, J., Lin, W., Schlesinger, F. et al. (2012), "Landscape of Transcription in Human Cells," *Nature*, 489, 101–108. [975]
- Flicek, P., Amode, M. R., Barrell, D., Beal, K., Brent, S., Chen, Y., Clapham, P., Coates, G., Fairley, S., Fitzgerald, S. et al. (2011), "Ensembl 2011," *Nucleic Acids Research*, 39, D800. ISSN 0305-1048. [978]
- Glaus, P., Honkela, A., and Rattray, M. (2012), "Identifying Differentially Expressed Transcripts From RNA-Seq Data With Biological Variation," *Bioinformatics*, 28, 1721–1728. [976]
- Honer, W. G., Falkai, P., Bayer, T. A., Xie, J., Hu, L., Li, H. Y., Arango, V., Mann, J. J., Dwork, A. J., and Trimble, W. S. (2002), "Abnormalities of Snare Mechanism Proteins in Anterior Frontal Cortex in Severe Mental Illness," *Cerebral Cortex*, 12, 349–356. [983]
- Huang da, W., Sherman, B. T., Lempicki, R. A. (2009), "Systematic and Integrative Analysis of Large Gene Lists Using David Bioinformatics Resources," *Nature Protocols*, 4, 44–57. [983]
- Hunter, J. E., Schmidt, F. L., and Jackson, G. B. (1982), *Meta-Analysis*, Beverly Hills, CA: Sage Publications. [985]
- Jiang, H. and Salzman, J. (2012), "Statistical Properties of an Early Stopping Rule for Resampling-Based Multiple Testing," *Biometrika*, 99, 973–980. [979]
- Jiang, H., and Wong, W. H. (2009), "Statistical Inferences for Isoform Expression in RNA-Seq," *Bioinformatics*, 25, 1026–1032. [975]
- Katz, Y., Wang, E. T., Airoldi, E. M., and Burge, C. B. (2010), "Analysis and Design of RNA Sequencing Experiments for Identifying Isoform Regulation," *Nature Methods*, 7, 1009–1015. [975]
- Keane, T. M., Goodstadt, L., Danecek, P., White, M. A., Wong, K., Yalcin, B., Heger, A., Agam, A., Slater, G., Goodson, M. et al. (2011), "Mouse Genomic Variation and Its Effect on Phenotypes and Gene Regulation," *Nature*, 477, 289–294. [983]
- Langmead, B., Hansen, K. D., and Leek, J. T. (2010), "Cloud-Scale RNA-Sequencing Differential Expression Analysis With Myrna," *Genome Biology*, 11, R83. [977]
- Leng, N., Dawson, J. A., Thomson, J. A., Ruotti, V., Rissman, A. I., Smits, B. M. G., Haag, J. D., Gould, M. N., Stewart, R. M., and Kendziorski, C. (2013), "EBSeq: An Empirical Bayes Hierarchical Model for Inference in RNA-Seq Experiments," *Bioinformatics*, 29, 1035–1043. [976]
- Li, B., Ruotti, V., Stewart, R. M., Thomson, J. A., and Dewey, C. N. (2010), "RNA-Seq Gene Expression Estimation With Read Mapping Uncertainty," *Bioinformatics*, 26, 493–500. [975]
- Li, J. J., Jiang, C. R., Hu, Y., Brown, B. J., Huang, H., and Bickel, P. J. (2011a), "Sparse Linear Modeling of RNA-Seq Data for Isoform Discovery and Abundance Estimation," *Proceedings of the National Academy of Sciences of the United States of America*, 108, 19867–19872. [975,977]
- Li, W., Feng, J., and Jiang, T. (2011b), "Isolasso: A Lasso Regression Approach to RNA-seq Based Transcriptome Assembly," *Journal of Computational Biology*, 18, 1693–1707. [975,977]
- Mazumder, R., Friedman, J. H., and Hastie, T. (2011), "Sparsenet: Coordinate Descent With Nonconvex Penalties," *Journal of the American Statistical Association*, 106, 1125–1138. [978]
- Mortazavi, A., Williams, B. A., McCue, K., Schaeffer, L., and Wold, B. (2008), "Mapping and Quantifying Mammalian Transcriptomes by RNA-Seq," *Nature Methods*, 5, 621–628. [975]
- Müller, D. J., Klempner, T. A., De Luca, V., Sicard, T., Volavka, J., Czobor, P., Sheitman, B. B., Lindenmayer, J. P., Citrome, L., McEvoy, J. P., Lieberman, J. A., Honer, W. G., and Kennedy, J. L. (2005), "The Snap-25 Gene May be Associated With Clinical Response and Weight Gain in Antipsychotic Treatment of Schizophrenia," *Neuroscience Letters*, 379, 81–89. [983]
- Pachter, L. (2011), "Models for Transcript Quantification From RNA-Seq," *Arxiv preprint arXiv:1104.3889*. [975]
- Pan, Q., Shai, O., Lee, L. J., Frey, B. J., and Blencowe, B. J. (2008), "Deep Surveying of Alternative Splicing Complexity in the Human Transcrip-

- tome by High-Throughput Sequencing," *Nature Genetics*, 40, 1413–1415. [975]
- Purdum, E., Simpson, K. M., Robinson, M. D., Conboy, J. G., Lapuk, A. V., and Speed, T. P. (2008), "Firma: A Method for Detection of Alternative Splicing From Exon Array Data," *Bioinformatics*, 24, 1707–1714. [975]
- Rashid, N. U., Giresi, P. G., Ibrahim, J. G., Sun, W., and Lieb, J. D. (2011), "Zinba Integrates Local Covariates With Dna-Seq Data to Identify Broad and Narrow Regions of Enrichment, Even Within Amplified Genomic Regions," *Genome Biology*, 12, R67. [979]
- Richard, H., Schulz, M. H., Sultan, M., Nürnberger, A., Schrunner, S., Balzereit, D., Dagand, E., Rasche, A., Lehrach, H., Vingron, M., Haas, S. A., and Yaspo, M. L. (2010), "Prediction of Alternative Isoforms From Exon Expression Levels in RNA-Seq Experiments," *Nucleic Acids Research*, 38, e112–e112. [975]
- Roberts, A., Trapnell, C., Donaghey, J., Rinn, J. L., Pachter, L. (2011), "Improving RNA-Seq Expression Estimates by Correcting for Fragment Bias," *Genome Biology*, 12, R22. [975]
- 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. [977]
- Salzman, J., Jiang, H., and Wong, W. H. (2011), "Statistical Modeling of RNA-Seq Data," *Statistical Science*, 26, 62–83. [975]
- Sommer, J. U., Schmitt, A., Heck, M., Schaeffer, E. L., Fendt, M., Zink, M., Nieselt, K., Symons, S., Petroianu, G., Lex, A., Herrera-Marschitz, M., Spanagel, R., Falkai, P., and Gebicke-Haerter, P. J. (2010), "Differential Expression of Presynaptic Genes in a Rat Model of Postnatal Hypoxia: Relevance to Schizophrenia," *European Archives of Psychiatry and Clinical Neuroscience*, 260, 81–89. [983]
- Sun, W. (2012), "A Statistical Framework for Eql Mapping Using RNA-Seq Data," *Biometrics*, 68, 1–11. [978]
- Sun, W., Ibrahim, J. G., and Zou, F. (2010), "Genomewide Multiple-Loci Mapping in Experimental Crosses by Iterative Adaptive Penalized Regression," *Genetics*, 185, 349–359. [977,978]
- Thompson, P. M., Sower, A. C., and Perrone-Bizzozero, N. I. (1998), "Altered Levels of the Synaptosomal Associated Protein Snap-25 in Schizophrenia," *Biological Psychiatry*, 43, 239–243. [983]
- Thorisson, G. A., Smith, A. V., Krishnan, L., and Stein, L. D. (2005), "The International Hapmap Project Web Site," *Genome Research*, 15, 1592–1593. [980]
- Tibshirani, R. (1996), "Regression Shrinkage and Selection via the Lasso," *Journal of the Royal Statistical Society, Series B*, 58, 267–288. [977]
- Trapnell, C., Pachter, L., and Salzberg, S. L. (2009), "TopHat: Discovering Splice Junctions With RNA-Seq," *Bioinformatics*, 25, 1105–1111. [976,979,983]
- Trapnell, C., Williams, B. A., Pertea, G., Mortazavi, A., Kwan, G., van Baren, M. J., Salzberg, S. L., Wold, B. J., and Pachter, L. (2010), "Transcript Assembly and Quantification by RNA-Seq Reveals Unannotated Transcripts and Isoform Switching During Cell Differentiation," *Nature Biotechnology*, 28, 511–515. [975,980]
- Trapnell, C., Hendrickson, D. G., Sauvageau, M., Goff, L., Rinn, J. L., and Pachter, L. (2013), "Differential Analysis of Gene Regulation at Transcript Resolution With RNA-Seq," *Nature Biotechnology*, 31, 46–53. [976,980]
- Wang, E. T., Sandberg, R., Luo, S., Khrebtkova, I., Zhang, L., Mayr, C., Kingsmore, S. F., Schroth, G. P., and Burge, C. B. (2008), "Alternative Isoform Regulation in Human Tissue Transcriptomes," *Nature*, 456, 470–476. ISSN 0028-0836. [975]
- Wang, G. S., and Cooper, T. A. (2007), "Splicing in Disease: Disruption of the Splicing Code and the Decoding Machinery," *Nature Reviews Genetics*, 8, 749–761. [975]
- Wang, Z., Gerstein, M., and Snyder, M. (2009), "RNA-Seq: A Revolutionary Tool for Transcriptomics," *Nature Reviews Genetics*, 10, 57–63. [975]
- Xia, Z., Wen, J., Chang, C. C., and Zhou, X. (2011), "Nsmmap: A Method for Spliced Isoforms Identification and Quantification from RNA-seq," *BMC Bioinformatics*, 12, 162. [975,977]
- Xing, Y., Yu, T., Wu, Y. N., Roy, M., Kim, J., and Lee, C. (2006), "An Expectation-Maximization Algorithm for Probabilistic Reconstructions of Full-Length Isoforms From Splice Graphs," *Nucleic Acids Research*, 34, 3150–3160. [975]
- Zhao, P., and Yu, B. (2006), "On Model Selection Consistency of Lasso," *The Journal of Machine Learning Research*, 7, 2541–2563. [978]
- Zou, H. (2006), "The Adaptive Lasso and Its Oracle Properties," *Journal of the American Statistical Association*, 101, 1418–1429. [977,978]