# A novel Bayesian approach to reconstructing and quantifying transcripts from RNA sequencing data

October 2, 2015

**Abstract**

**Background**

High throughput sequencing methods for RNA (RNA-Seq) have enabled the study of transcript structure and expression levels at unprecedented resolution. However accurate transcript reconstruction from RNA-Seq reads remains challenging, and analyses often summarize data by simply counting the reads mapping to each gene or exon.

We propose a fully Bayesian approach to reconstruct transcripts and quantify their abundances from RNA-Seq data. A key element of our approach is a novel generative statistical model for transcripts, which forms the prior distribution in our analyses. This model uses the concept of a "flexible exon" (flexon), which is an exon with multiple potential start and end locations. We combine this prior model for transcripts, with a prior on expression values, and with the RNA-Seq data, to infer the posterior distribution on transcripts and their expression levels.

**Results**

We assess the performance of our method, `altra`, both by comparing the fitted model with the raw RNA-Seq data, and by assessing its ability to identify eQTLs affecting splicing that were previously identified by an exon-level analysis. When we compare `altra` to a widely used transcript reconstruction method, Cufflinks, we find that, even

though our fitted model appears to capture many of the key features evident in the raw data, neither our method nor Cufflinks identifies the majority of the splice eQTLs with `altra` identifying 30%, and Cufflinks identifying 12% of the splice eQTLs (both at $p$-value$<$ 0.01).

**Conclusions**

Our analyses suggest that transcript-reconstruction-based methods in splice eQTL analyses should be complemented by exon-level analyses and confirms that, although promising, accurate transcript reconstruction from RNA-Seq data remains a challenging problem. We also outline a potential alternative approach to this problem that is suggested by our work.

**Key words:** RNA sequencing, transcript reconstruction, alternative transcripts, gene expression, splice eQTLs, multiple samples, Bayesian

# 1 Background

The use of RNA sequencing to measure gene expression has already provided new insights into transcription and splicing (e.g. Pan et al. (2008); Wang et al. (2008); Mortazavi et al. (2008); Blekhman et al. (2010); Pickrell et al. (2010b); Gonzàlez-Porta et al. (2012); Djebali et al. (2012); Barbosa-Morais et al. (2012); Merkin et al. (2012); Reyes et al. (2013); Ergun et al. (2013)), and how these processes are affected by genetic variation (e.g. Pickrell et al. (2010a); Montgomery et al. (2010); Sun and Hu (2013); Zhao et al. (2013); Lappalainen et al. (2013); Kurmangaliyev et al. (2013)). However a full understanding of the relationship between genetic variants and alternative transcription is hindered by the difficulty of accurately reconstructing and quantifying transcripts. These remain challenging because of both data limitations and transcriptome combinatorial complexity. As a result, despite numerous methods for reconstructing transcripts from RNA sequencing data (e.g. Denoeud et al.

(2008); Guttman et al. (2010); Filichkin et al. (2010); Montgomery et al. (2010); Trapnell et al. (2010); Zhao et al. (2011); Li et al. (2011a,b); Feng et al. (2011); Xia et al. (2011); Mezlini et al. (2012); Hiller and Wong (2013); Mangul et al. (2012); Lin et al. (2012); Behr et al. (2013); Bernard et al. (2013)), in practice expression analyses often rely on summarized expression at the level of annotated genes (e.g. Anders and Huber (2010); Robinson et al. (2010); Turro et al. (2011)), annotated or novel exons (e.g. Anders et al. (2012)) or events (e.g. Katz et al. (2010); Brooks et al. (2011); Seok et al. (2012); Barbosa-Morais et al. (2012); Merkin et al. (2012); Reyes et al. (2013); Anders et al. (2013)) without trying to reconstruct transcripts.

Here, we present novel methods for transcript reconstruction and differential expression estimation at the transcript level, using RNA sequencing data. The methods require a genome for the mapping of the reads (i.e. they are "genome-guided"), but can operate either with or without a gene annotation. Furthermore, the methods are designed to use partial gene annotation data – specifically, a list of candidate 3' and 5' splice sites – which are much easier to reliably obtain than full transcript annotations. Given a list of candidate splice sites (plus start and end sites), we specify a prior distribution for the set of possible transcripts, using the idea of a "flexible exon" ("flexon"), which captures the idea that transcripts may contain different "versions" of an exon, as well as contain different exons. By combining this prior distribution with a likelihood for the RNA-Seq data we perform Bayesian inference for the gene model, and the expression values for sequenced samples. As far as we are aware our method is the first fully Bayesian approach at this problem, and is one of the few methods that attempt to jointly reconstruct transcripts and estimate transcript expression values for multiple samples (see also Cliique Lin et al. (2012) and miTie Behr et al. (2013)). A software package `altra` implementing our method is available at https://github.com/esterpantaleo/altra.

We assess the performance of our method both by comparing the fitted model with

the raw RNA-Seq data, and by assessing its ability to identify splice eQTLs that were previously identified by an exon-level analysis. In the second analysis we compare with a widely-used method, Cufflinks Trapnell et al. (2010). We find that, even though our fitted model appears to capture many of the key features evident in the raw data, accurate transcript reconstruction and expression level estimation remains a challenging problem: neither method identifies the majority of the splice eQTLs in our assessment, with `altra` identifying 30%, and Cufflinks identifying 12% of the splice eQTLs (both at $p$-value$< 0.01$).

# 2 Methods

Our Bayesian approach has three key components: the likelihood for the RNA-Seq data, a prior distribution on the gene model (set of transcripts), and a hierarchical model on expression levels that combines information across individuals. We describe each of these in turn.

### The likelihood

Our likelihood is based on the likelihood in Hiller and Wong (2013), but we introduce an additional error term ($\epsilon$) to allow for mis-mapped reads, low-expressed transcripts and/or rare intronic reads. In addition we deal with multiple samples simultaneously: we assume that at each locus (gene) the $N$ samples share a common (unknown) set of $T$ transcripts $\tau$ – *the gene model* – with (unknown) expression values $\boldsymbol{\lambda}$, and derive a log-likelihood $l(\tau, \boldsymbol{\lambda}, \epsilon)$.

Let $\rho$ denote the set of observable (single-end) reads at the locus. Note that $\rho$ includes both spliced and non-spliced reads, but excludes reads that could not be observed due to overhang constraints Katz et al. (2010). The data $x_r^i$ are the counts of the number of reads of type $r \in \rho$ in sample $i$ ($i = 1, \ldots, N$). (Some counts may be 0, since some potentially observable reads will not occur in the data). We assume, as in Hiller and Wong (2013) that

the counts are independent and Poisson distributed:

$$x_r^i \mid \tau, \boldsymbol{\lambda}_i, \epsilon \sim Pois\left( (\sum_t z_{rt}(\tau)\lambda_t^i + \epsilon)C^i \right). \quad (1)$$

Here $\lambda_t^i$ denotes the expression level of transcript $t$ in sample $i$, $z_{rt}(\tau)$ is a binary indicator for whether transcript $t$ in the gene model $\tau$ could produce a read of type $r \in \rho$, and $C^i$ is a known normalization constant (we use the total number of reads mapped in sample $i$ divided by the average number of reads across samples, but more complex normalization procedures might be necessary, depending on the data Dillies et al. (2012)).

The independence assumption means that the total likelihood is obtained by multiplying (1) across samples $i$ and read types $r$. As in Salzman et al. (2010), this product can be simplified by grouping together terms corresponding to reads $r$ that are compatible with the same subset of transcripts. That is, reads $r$ that share the same "compatibility vector" $z_{r\cdot}$. Specifically, the overall log-likelihood based on the entire data $\boldsymbol{x}$ simplifies to

$$\begin{aligned} l(\tau, \boldsymbol{\lambda}, \epsilon; \boldsymbol{x}) &= \sum_i \sum_{\boldsymbol{z} \in \mathcal{Z}} \{x_{\boldsymbol{z}}^i \ \log\left[\left(\sum_t z_t \lambda_t^i + \epsilon\right) C^i\right]\} - \sum_i C^i \left(\sum_t l_t \lambda_t^i\right) \\ &\quad - \epsilon L \sum_i C^i + const \end{aligned}$$

where $x_{\boldsymbol{z}}^i$ denotes the number of reads in sample $i$ with compatibility vector $\boldsymbol{z}$, $\mathcal{Z}$ is the set of all possible compatibility vectors $\boldsymbol{z}$ (a set with cardinality $2^T$), $l_t = \sum_{r \in \rho} z_{rt}$ is the effective length of transcript $t$ (the length of the transcript, ignoring unmappable bases due to overhang constraints), and $L$ is the length of the locus (the number of bases in the locus).

## Priors on the gene model $\tau$

We assume that the number of transcripts, $T$, in $\tau$ is fixed by the user. Given $T$ we specify a prior on the gene model in 3 steps (see Figure 1).

1. We identify sets $S^{(3)}$ and $S^{(5)}$ of 3' and 5' splice sites that could occur at the locus. These can come from a combination of gene annotation databases, observed splice junctions in the data, and putative splice sites found by a segmentation method Nowak et al. (2011). (Note that in the latter two cases our prior is technically "data-dependent", but the use of the data here is relatively weak, serving only to reduce the support of the prior to regions that are consistent with the data.)

2. From these splice sites, we define "flexible exons" ("flexons"). Just as an exon can be defined by a single 3' and 5' splice site, a flexon is defined by a set $S_f^{(3)}$ of one or more 3' splice sites, and a set $S_f^{(5)}$ of one or more 5' splice sites, together with probability vectors ($\boldsymbol{p}_f^{(3)}$, $\boldsymbol{p}_f^{(5)}$) that specifiy how likely the flexon is to "use" each splice site. In other words, a flexon is a distribution on exons, and a random exon could be generated from this distribution by selecting a 3' splice site at random from $S_f^{(3)}$ according to $\boldsymbol{p}_f^{(3)}$ and a 5' splice site from $S_f^{(5)}$ according to $\boldsymbol{p}_f^{(5)}$. The number of flexons $F$, and the sets $S_f^{(3)}, S_f^{(5)}$ $(f = 1, \dots, F)$ for each flexon are computed deterministically from the sets $S^{(3)}$ and $S^{(5)}$ identified above [see Additional file 1]. The probability vectors $\boldsymbol{p}_f^{(3)}, \boldsymbol{p}_f^{(5)}$ are parameters of the prior, and assigned uniform prior distributions.

3. Finally, we assume that each transcript includes or excludes flexon $f$ independently, with probability $p_f$ $(f = 1, \dots, F)$, and assign a uniform prior to each $p_f$ [see Additional file 1].

To further illustrate the idea of a flexon, suppose that $S_f^{(3)}$ contains only one element and $S_f^{(5)}$ contains two elements. Then flexon $f$ will have two possible forms that share the same 3' splice site and have alternative 5' splice sites, one upstream of the other. If the probabilities associated with $f$ are $p_f = 1/2$, and $\boldsymbol{p}_f^{(5)} = (1/3, 2/3)$ then one out of two transcripts will include the flexon and, if the flexon is included, the longer form will be twice as frequent as the short form.

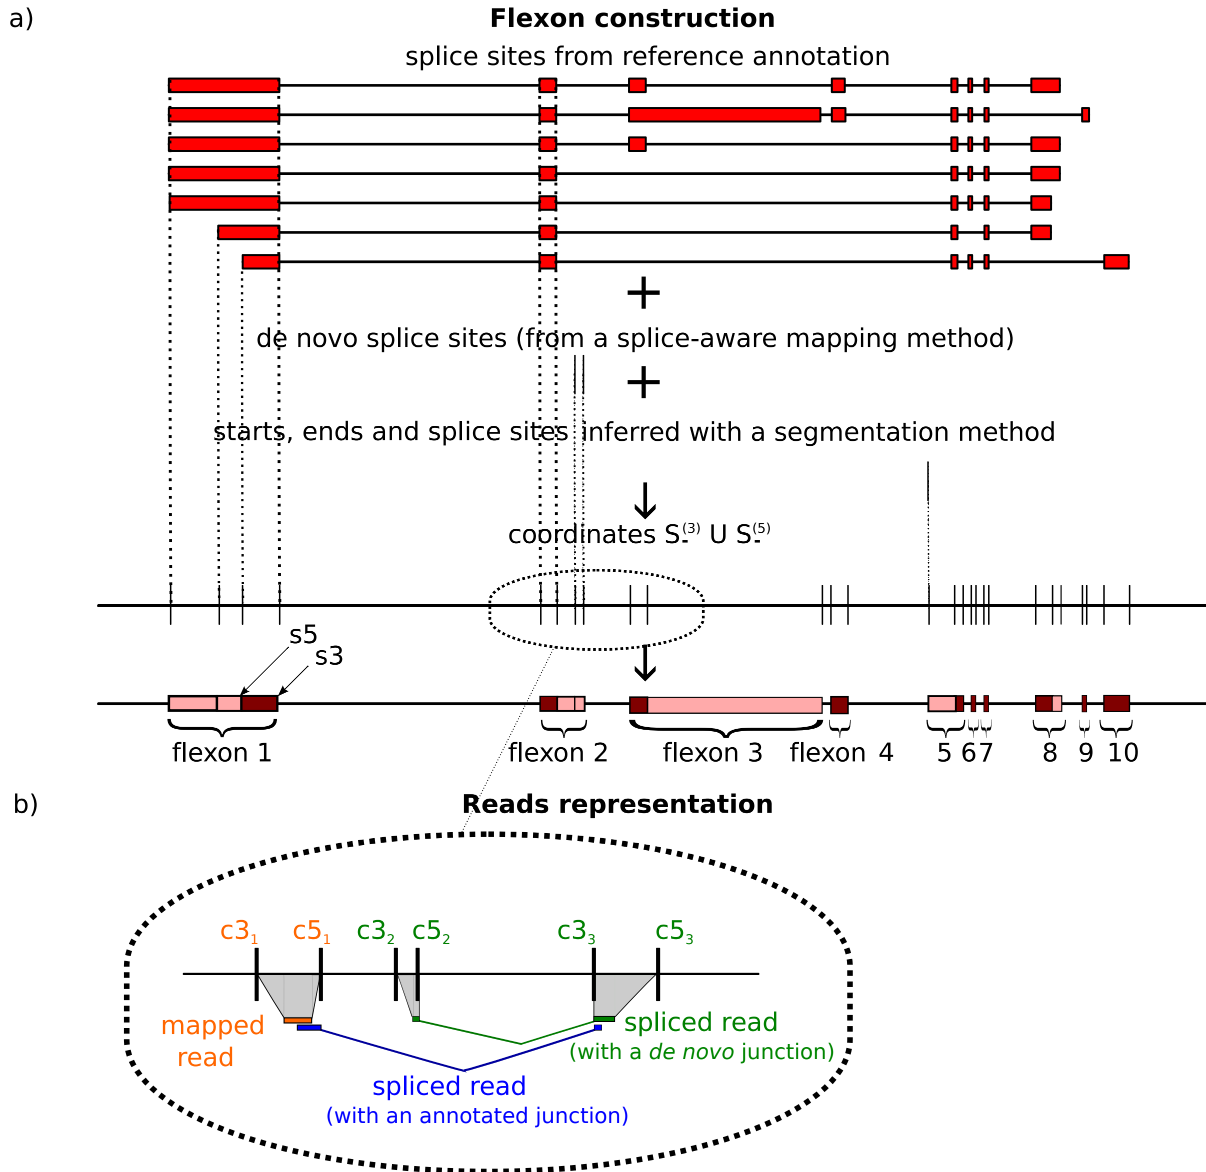


Figure 1: **Flexon construction in gene *HMGN1* (reverse strand).** First we construct the sets of negative 3' and 5' splice sites ($S_{-}^{(3)}$ and $S_{-}^{(5)}$) from the RefSeq, UCSC and Ensembl (hg18) annotations, and/or from a splice-aware mapping method (TopHat Trapnell et al. (2009)) and/or from a segmentation method (FLLatNowak et al. (2011)). Then we identify segments $[s3, s5]$ with $s3 \in S_{-}^{(3)}$ and $s5 \in S_{-}^{(5)}$ (dark red in Figure; s3 and s5 in flexon 1 are indicated by arrows). Finally we collect all 3' and 5' splice sites around $[s3, s5]$ that could generate an exon larger than the minimum exon length and smaller than the maximum exon length (these lengths are parameters of the model). For illustrative purposes, here we are representing only a subset of the possible flexon configurations; for example, we represent three different forms of flexon 1 that share the same 3' splice site s3 and use different 5' splice sites. b) **Representation of three different reads that map to gene *HMGN1*.** Reads are represented by coordinates in $S_{-}^{(3)}$7and $S_{-}^{(5)}$. The orange read is represented by coordinates $(c3_1, c5_1)$ and the blue read (a spliced read) by coordinates $(c3_1, c5_1, c3_3, c5_3)$. [See Additional file 1 for more details.]

The fact that $p_f$, $\boldsymbol{p}_f^{(3)}$, and $\boldsymbol{p}_f^{(5)}$ are assigned prior distributions and then inferred from the data [see Additional file 1] allows for sharing of information across transcripts. For example, some flexons might be "constitutive flexons" and be used by most or all transcripts (resulting in high estimated $p_f$). And, for each flexon, some alternative 3' (5') splice sites might be more commonly used than others, which would be reflected in the estimates of $p_f^{(3)}$ ($p_f^{(5)}$). Indeed, many flexons may have effectively no variation in splice site usage, in which case they reduce to the regular concept of an exon.

Some motivation for this model can be gained by examining the simple gene model in Figure 2f. Many exons occur in every transcript with the same 3' and 5' splice sites - these would be captured in our model by flexons that are always included ($p_f = 1$) and have no variation in splice site usage. However, the exon numbered 4 occurs in all transcripts, but has a slightly modified form (different 3' splice site) in one transcript - this would also be captured by a flexon that occurs in all transcripts ($p_f = 1$) but with variation in splice site usage. Finally there is an exon (labelled 1) that occurs in only some transcripts, and has multiple forms, which again can be captured by our flexon-based model.

## Hierarchical model for expression levels $\boldsymbol{\lambda}$

As is conventional, we model expression levels $\boldsymbol{\lambda}$ on a log scale. Specifically we use a hierarchical model on the log-expression values $\theta_t^i = \log \lambda_t^i$ to share information across samples:

$$\theta_t^i | \bar{\theta}_t, \sigma_{\theta,t}^2 \sim \mathcal{N}(\bar{\theta}_t, \sigma_{\theta,t}^2) \qquad \text{for } t \in \{1, \ldots, T\} \text{ and } i \in \{1, \ldots, N\} \tag{2}$$

with

$$\bar{\theta}_t | \sigma_{\theta,t}^2, c \sim \mathcal{N}(0, c\, \sigma_{\theta,t}^2) \qquad \text{for } t \in \{1, \ldots, T\} \tag{3}$$

$$\sigma_{\theta,t}^2 | \alpha, \beta \sim \mathcal{IG}(\alpha, \beta), \tag{4}$$

where $\alpha = 4$, $\beta = 1$ and $c = 2$. Here $\bar{\theta}_t$ represents the mean (log-) expression of transcript $t$ across samples. Assuming that the samples share a common mean helps to combine information among them.

Similarly, if $\delta = \log \epsilon$, we assume

$$\delta | \mu_\delta, \sigma_\delta^2 \sim \mathcal{N}(\mu_\delta, \sigma_\delta^2).$$

where $\mu_\delta = -4$ and $\sigma_\delta = 1$.

We checked that results were robust to different values of $\alpha$, $\beta$, $c$, $\mu_\delta$ and $\sigma_\delta$.

## Exploring the posterior distribution

We use a Metropolis Hastings (MH) algorithm to explore the posterior distribution $p(\tau, \boldsymbol{\theta}, \delta | \boldsymbol{x})$ given the data $\boldsymbol{x}$ and the above priors on $\tau, \boldsymbol{\theta}$, and $\delta$.

We initialize the MH algorithm by sampling $\boldsymbol{\theta}$ and $\delta$ from the prior and by randomly sampling $T$ transcripts. If annotation is available, we choose a random set of $T$ annotated transcripts as a sensible start.

After the initialization, we update $\theta_t^i$ and $\delta$ with an independent random walk Metropolis algorithm given the other variables. To update $\tau$, we select a random transcript and we propose one or a combination of several different types of move: including/excluding one of its flexons, changing the 3' or 5' splice site used at a flexon, or more complex moves that involve pairs of transcripts - like swapping exons between transcripts, or recombining transcripts. [See Additional file 1 for a detailed description.]

# Results

## Illustrative examples

We begin with a simple example to illustrate how combining information across multiple samples can improve transcript reconstruction accuracy. Our example involves two transcripts expressed in two samples A and B (Figure 2). The transcripts differ only by inclusion/exclusion of one exon (exon labelled 3 in Figure). Exon skipping is the most common pattern of alternative splicing in humans Wang et al. (2008). The expression level of each transcript in each sample is shown in the figure (panel a): both transcripts are expressed in each sample, but in sample A only the transcript that skips the exon is highly expressed, and in sample B only the other transcript is highly expressed. We simulated reads for each sample according to a Poisson model (Figure 2b) and assumed them to be mapped without error.

We applied `altra` to these simulated data, both one sample at a time, and jointly to both samples. (In applying `altra` we used the full gene annotation, Figure 2f, to identify putative splice sites and transcription start and end sites.) The results show that, when using data from only one sample, we can accurately reconstruct only the highly expressed transcript (Figures 2c,d). However, analyzing the two samples jointly accurately reconstructs both transcripts (Figure 2e).

We now illustrate our method on real data. We consider two genes, *COMMD4* and *TOMM40L*, that are expressed in the Illumina Brain RNA-Seq dataset in Au et al. (2010). We chose these data because they were also used by Hiller and Wong (2013) to illustrate the performance of their method, Montebello, and so this allows us to compare our results with theirs (see Figures 4 and 5 in Hiller and Wong (2013)).

Following Hiller and Wong (2013), we mapped the reads (paired-end reads of length 50bp, from three technical replicates which we pooled into a single sample) to hg18 with SpliceMap

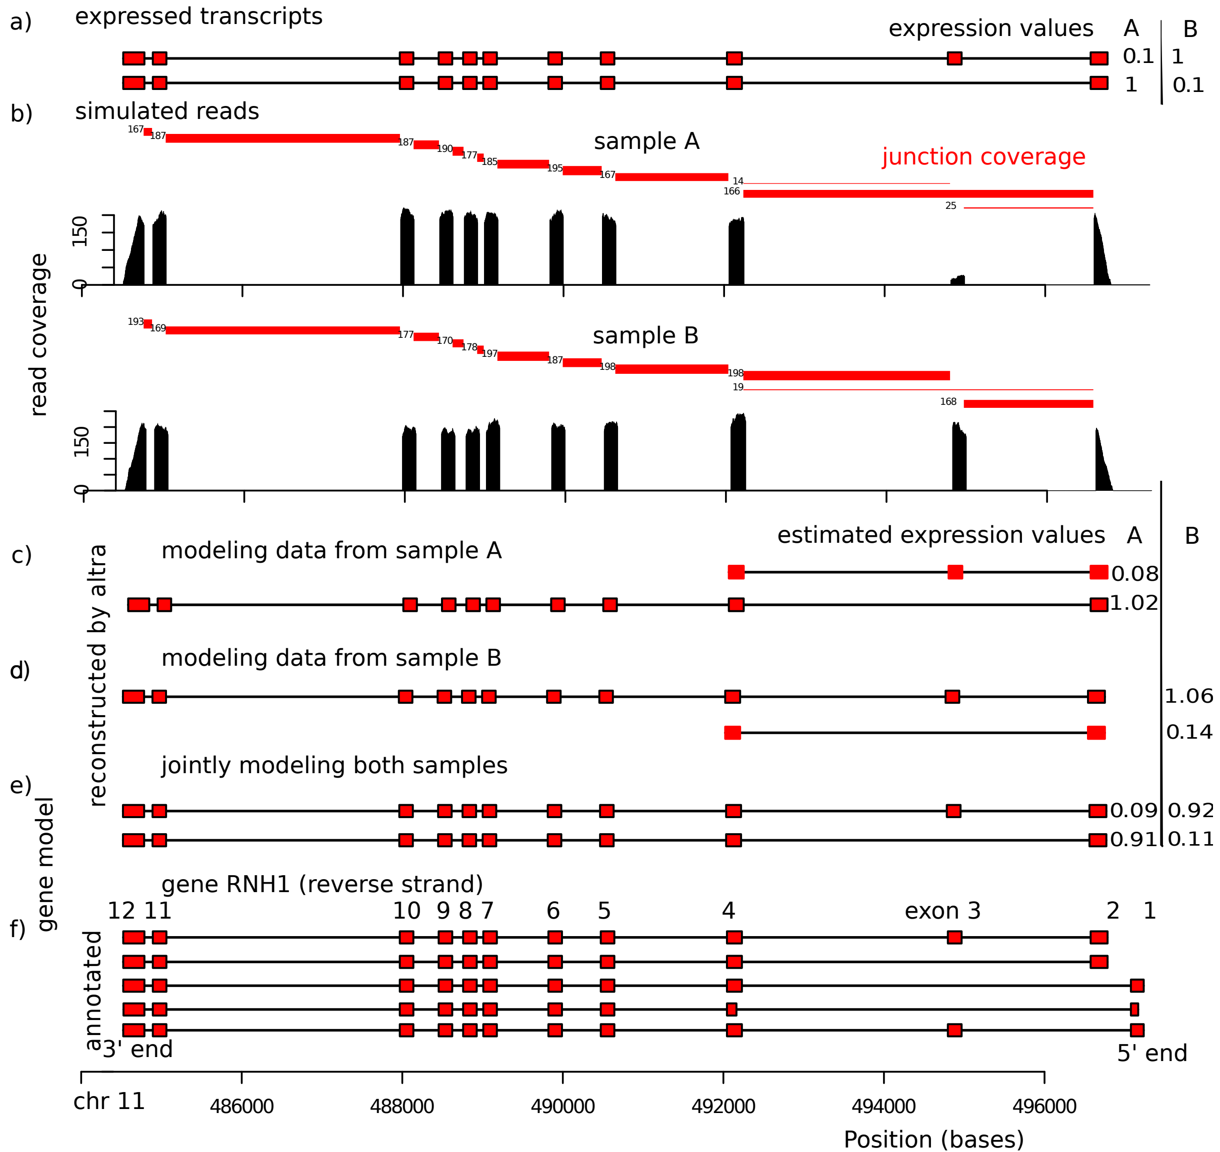


Figure 2: **Combining information across multiple samples can improve alternative transcript reconstruction accuracy.** a) Two annotated alternative transcripts of gene *RNH1*. Isoforms differ by inclusion/exclusion of a single exon. Simulated expression levels for each transcript in samples A and B are reported on the right. b) Reads (46bp) simulated from transcripts in (a). c) Gene model reconstructed by `altra` using only data from sample A. Expression values are reported on the right. d) Gene model reconstructed by `altra` using only data from sample B. Expression values are reported on the right. e) Gene model reconstructed by `altra` by combining information across samples A and B. Average expression values for sample A and sample B are reported on the right. f) Full Ensembl (4.30) gene model.

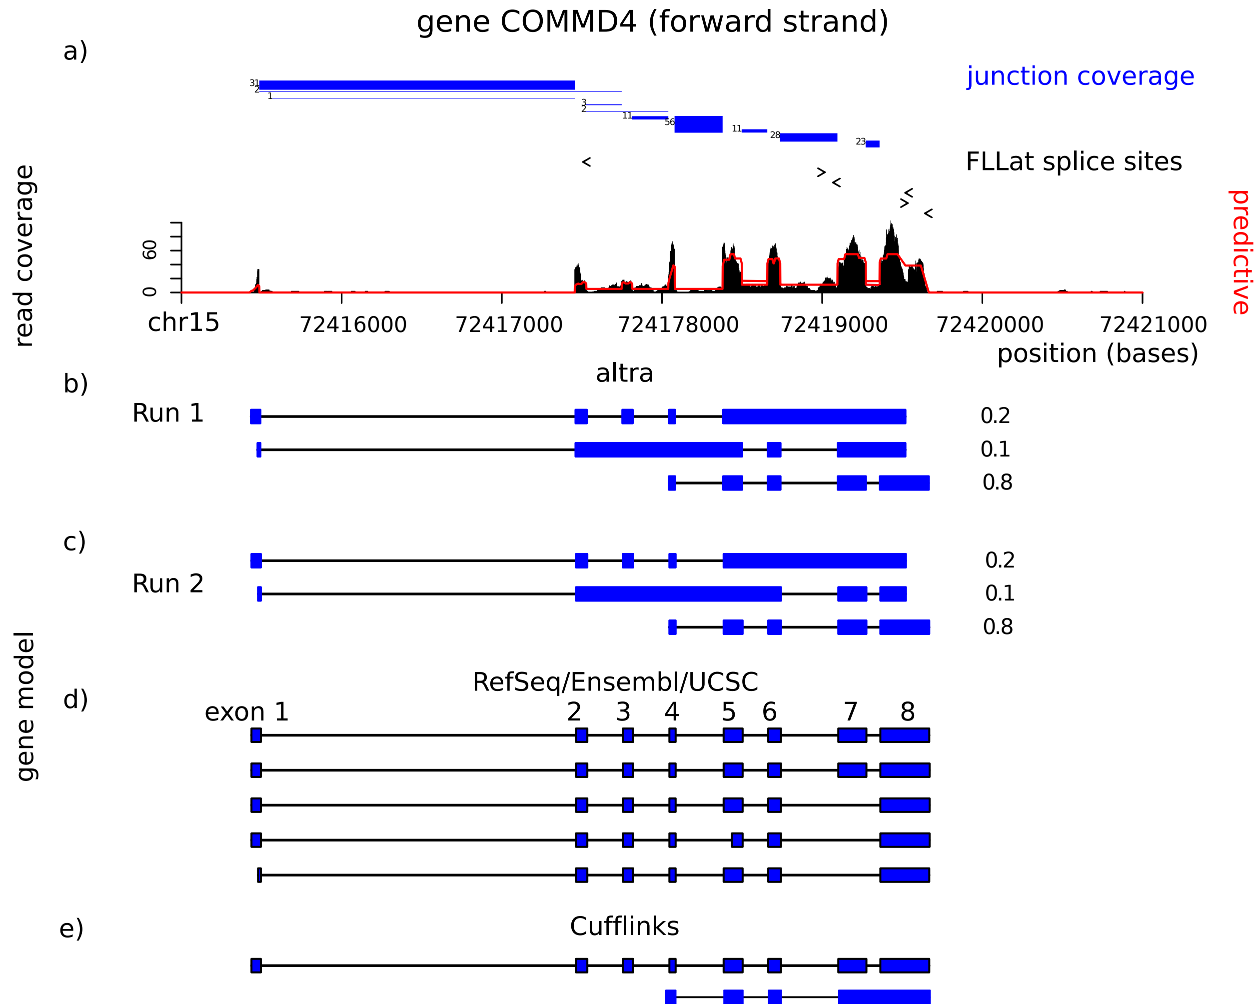


Figure 3: **Gene *COMMD4***, data from Illumina Brain RNA-Seq Au et al. (2010). a) Junction coverage: Each line represents a gap; the number of spliced reads that map to the genome with that gap is reported next to the gap (the ticker the line, the higher the number of spliced reads). FLLat splice sites: Splice sites identified by FLLat Nowak et al. (2011). In red: Predictive distribution inferred by `altra`. b) Multiple runs of `altra` can return different gene models. We report two different gene models with corresponding expression levels. (We ran `altra` with options -J 0 -M 1, i.e., without filtering lowly expressed junctions.) d) Reference annotation (hg18). e) Gene model reconstructed by Cufflinks (v2.0.2, which was the version available when these comparisons were performed). We ran Cufflinks on the full dataset with two different choices of the cutoff parameter F (0 and 0.2): if a transcript has relative expression less than F% Cufflinks will not report it. Cufflinks returned the same gene model for both choices of F.

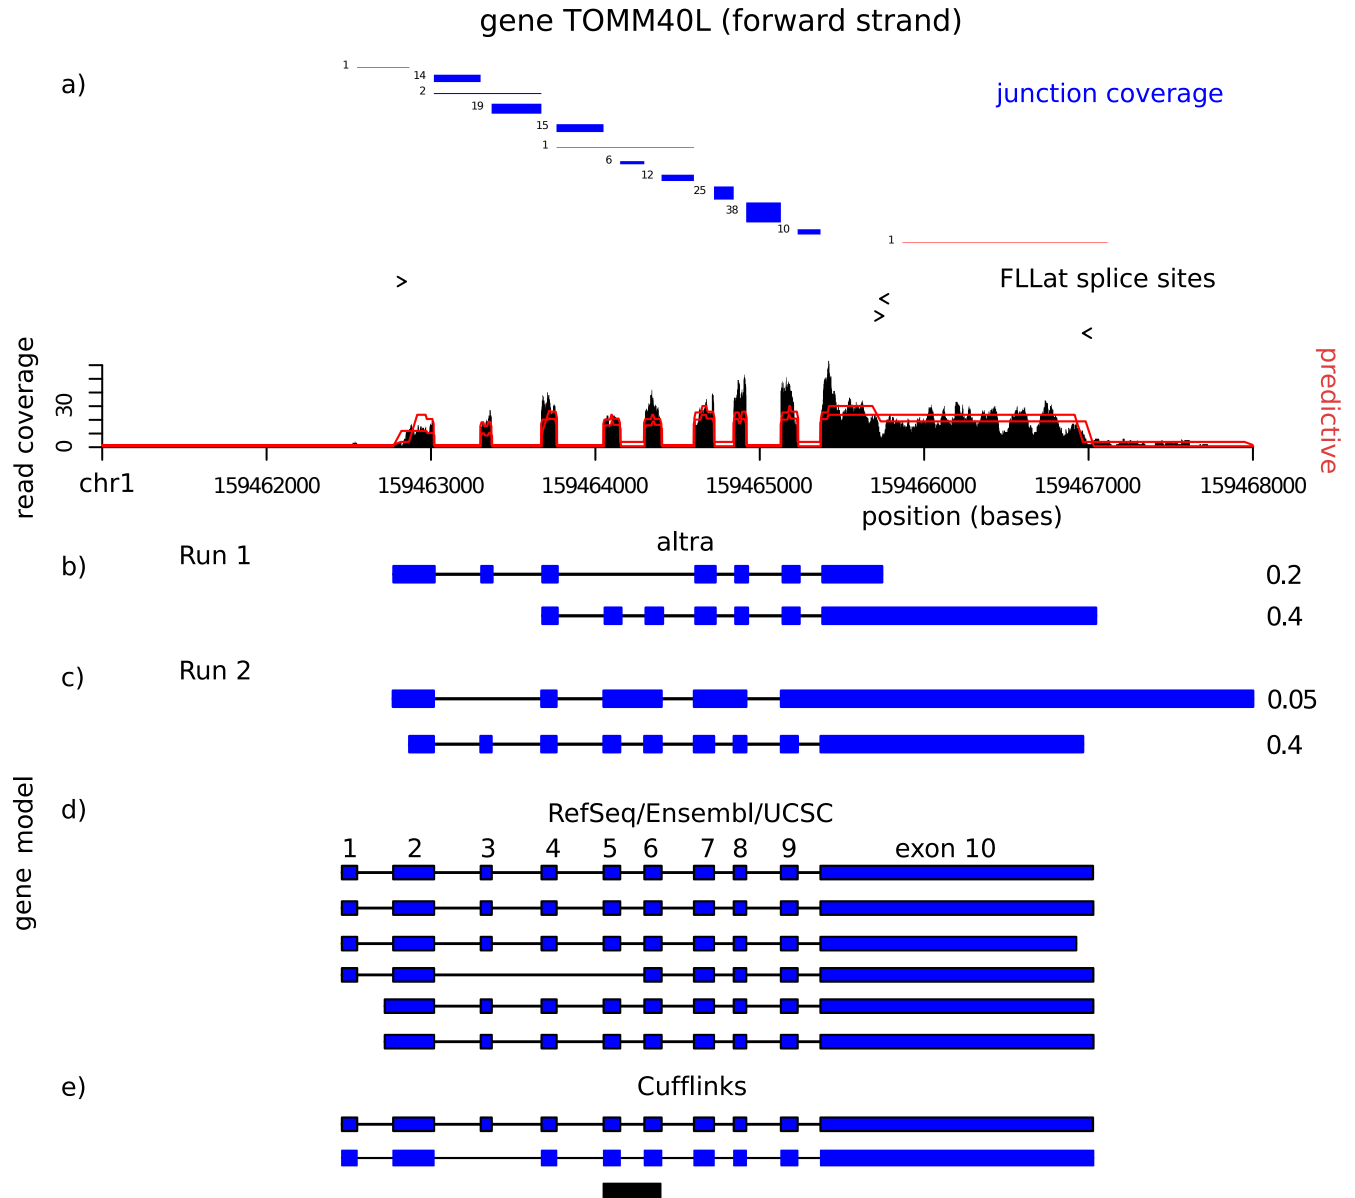


Figure 4: **Gene *TOMM40L***, data from Illumina Brain RNA-Seq Au et al. (2010). a) Junction coverage, splice sites identified by FLLat, read coverage and (in red) predictive distribution inferred by `altra`. b) Multiple runs of `altra` can return different gene models. We report two different gene models with the corresponding expression levels. (We ran `altra` with options -J 0 -M 1, i.e., without filtering lowly expressed junctions.) d) Reference annotation (hg18). e) Gene model reconstructed by Cufflinks (v2.0.2). Cufflinks returned the same gene model for both F=0 and F=0.2.

(v3.3.5.2) Au et al. (2010). Figures 5a and 6a show the resulting read and junction coverage across each gene. We applied `altra` to each locus, each time running `altra` multiple times to assess consistency of solutions across runs (which is a common way to assess the reliability of Monte-Carlo-based inference methods). After experimenting with different number of transcripts ($T = 2 - 5$) we chose to show illustrative results for $T = 3$ for *COMMD4* and $T = 2$ for *TOMM40L*. (Larger values of $T$ tended to produce larger discrepancies in inferred gene models, suggesting more extreme convergence problems.)

For each locus we found that different runs of `altra` could produce different inferred gene models (Figures 5b,c, 6b,c). To visually examine the fit of each gene model to the data we overlaid the fitted (predicted) coverage distribution on the observed coverage (Figures 5a, 6a). In each case we found that results from different runs on the same data, although differing in their inferred gene model, typically provided similar fitted coverage distributions. Furthermore, the fitted models captured many key features evident in the raw data. For example, in *COMMD4* (Figure 5b,c), `altra` infers two long exons in locations that are annotated as intronic but where we observe lots of expression. It also infers an unannotated transcription start site downstream of the annotated start site, which is consistent with the lower coverage in the data of the first three exons. Similarly, in *TOMM40L* (Figure 6b,c), `altra` infers a new start and a new end as well as novel splicing events that are supported by the coverage pattern. Thus, although the fact that different runs produce different results suggests that our Monte Carlo inference scheme is failing to fully converge, the comparison of the fitted models with the raw data suggests that convergence is good enough to find solutions that are largely consistent with the data.

We also ran Montebello on these data, and found that it also exhibited convergence problems, returning different gene models in different runs (data not shown). This is despite the fact that Montebello implements a parallel tempering approach to improve convergence. (We also implemented a parallel tempering approach in `altra`, but it did not yield sufficient

improvement to justify the additional computational burden.) To provide a specific set of inferred transcripts for comparison we take the published results from Figures 4 and 5 in Hiller and Wong (2013). Unsurprisingly given the complexity of the problem, their inferred gene models differ markedly from those inferred by `altra`. For example, at *COMMD4*, Montebello uses annotated ends while `altra` uses an alternative end site upstream of the annotated end. The expression pattern at the 5' end of *COMMD4* is compatible with both scenarios and we don't know which model is correct, if any. At *TOMM40L* the highest expressed transcript inferred by Montebello (their Figure 5) has a final long exon. Although again we do not know the true model, this particular inference seems inconsistent with the data (only a few reads map to that region). We note that the main conceptual difference between Montebello and `altra` is our novel prior distribution on transcript models that shares information across transcripts. Therefore, in principle differences in the results between the methods could help assess the gains that this novel prior provides. However in practice the problem is complex enough that observed differences could be due to many factors, including detailed implementational issues.

For comparison with these likelihood-based approaches, we also applied the graph-based approach Cufflinks Trapnell et al. (2010) to these data. The results are shown in Figures 5e and 6e. Cufflinks selects the number of transcripts automatically, and selected 2 transcripts for each locus (plus a single one-exon transcript at *TOMM40L* overlapping two annotated exon and an annotated intron). For both *COMMD4* and *TOMM40L* some of Cufflinks's inferred transcripts exactly match an annotated transcript; however Cufflinks uses the annotated transcripts when performing its inference (by adding simulated reads from the annotation to the observed reads Roberts et al. (2011)), and so the match between the inference and the annotation does not guarantee that the annotated transcripts are present in the data. At *COMMD4* a key qualitative difference between the results from `altra` and Cufflinks is that `altra` infers intron retention between multiple exons (2 and 3; 3 and 4; 5

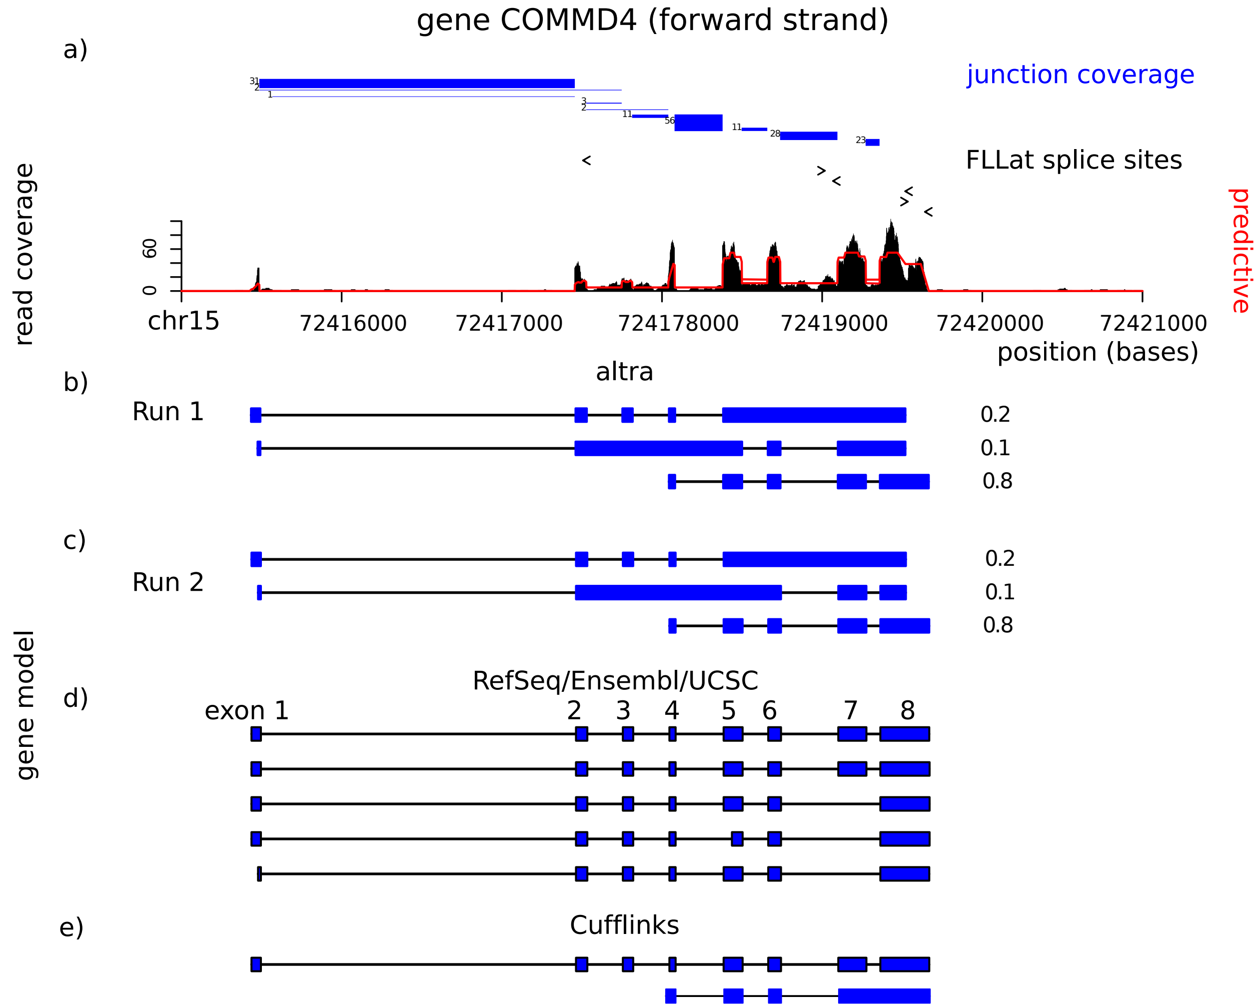


Figure 5: **Gene *COMMD4***, data from Illumina Brain RNA-Seq Au et al. (2010). a) Junction coverage: Each line represents a gap; the number of spliced reads that map to the genome with that gap is reported next to the gap (the ticker the line, the higher the number of spliced reads). FLLat splice sites: Splice sites identified by FLLat Nowak et al. (2011). In red: Predictive distribution inferred by `altra`. b) Multiple runs of `altra` can return different gene models. We report two different gene models with corresponding expression levels. (We ran `altra` with options -J 0 -M 1, i.e., without filtering lowly expressed junctions.) d) Reference annotation (hg18). e) Gene model reconstructed by Cufflinks (v2.0.2, which was the version available when these comparisons were performed). We ran Cufflinks on the full dataset with two different choices of the cutoff parameter F (0 and 0.2): if a transcript has relative expression less than F% Cufflinks will not report it. Cufflinks returned the same gene model for both choices of F.

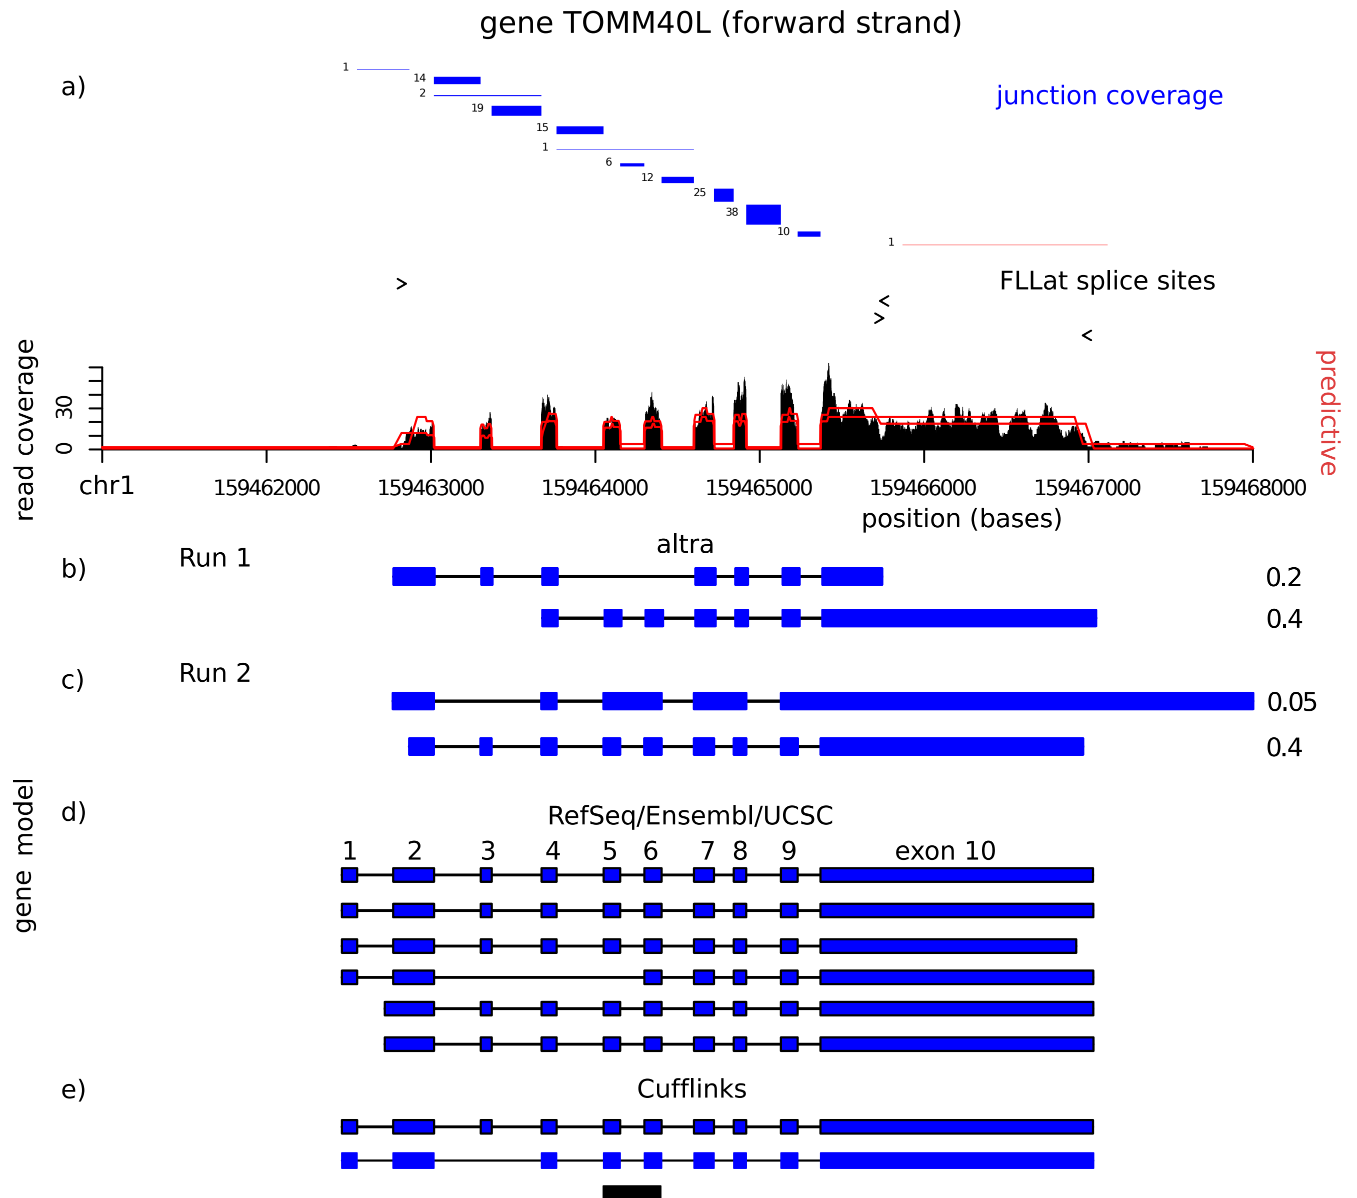


Figure 6: **Gene *TOMM40L***, data from Illumina Brain RNA-Seq Au et al. (2010). a) Junction coverage, splice sites identified by FLLat, read coverage and (in red) predictive distribution inferred by `altra`. b) Multiple runs of `altra` can return different gene models. We report two different gene models with the corresponding expression levels. (We ran `altra` with options -J 0 -M 1, i.e., without filtering lowly expressed junctions.) d) Reference annotation (hg18). e) Gene model reconstructed by Cufflinks (v2.0.2). Cufflinks returned the same gene model for both F=0 and F=0.2.

and 6; 6 and 7; 7 and 8) while Cufflinks infers intron retention only between exons 7 and 8. Again, one can see in the data the intronic reads that drive the inferences being made by `altra`, although we do not know for certain whether these inferences are correct.

These examples and results serve more to emphasize the complexity of this problem than to compare the accuracy of the different methods, which is difficult to assess. A simple comparison of observed and predicted coverage for each method is insufficient because methods might be capturing not only biological features (alternative promoter usage, alternative polyadenylation, alternative splicing) but also technical artifacts (e.g. positional biases, sequence-specific biases and mapping issues Griebel et al. (2012)). In fact, a 3' bias could explain why both the first exon in *TOMM40L* and the first three exons in *COMMD4* are lowly expressed compared to other exons and not alternative start site usage as predicted by `altra`. Analogously, mapping biases could explain the dip in expression in the last exon of *COMMD4* and *TOMM40L*, and not alternative end site usage as predicted by `altra`. Neither Montebello nor `altra` is modeling biases. An option can be selected for Cufflinks to model biases; however, as the authors mention in the manual, positional bias correction reduces accuracy on certain datasets in some genes. For this reason we did not use that option when running Cufflinks.

## Differential expression at the transcript level

Given the difficulty of directly assessing transcript reconstruction methods when the true gene model is unknown, we turned to an indirect way to assess the performance of `altra` on real data. We chose a dataset in which splice eQTLs (i.e., SNPs affecting splicing) have been previously identified using an *exon-level analysis* Pickrell et al. (2010a). This dataset consists of sequenced RNA (single-end reads of length 46bp) from 69 extensively genotyped human lymphoblastoid cell lines (LCLs) derived from unrelated Nigerian individuals. We used `altra` and Cufflinks to reconstruct gene models in these loci, and `altra` to estimate the expression

level of each transcript in each individual. We then tested the splice eQTLs for association with the reconstructed transcript expression levels. If reconstruction is accurate, we expect to have greater power to detect the splice eQTL than if reconstruction is inaccurate.

Pickrell et al. Pickrell et al. (2010a) searched for splice eQTLs by treating the fraction of reads mapped to each exon (of all the reads in the gene) as a quantitative trait. Specifically, they performed simple linear regressions of these fractions (suitably normalized, and controlling for unmeasured confounders) against all SNPs within 200 kb of the gene. In this way they identified 187 associations between expression of an exon and a SNP (FDR =10%).

We mapped the reads to hg19 with Tophat (v2.0.0) Trapnell et al. (2009) and ran `altra` in two ways: first, taking account of which reads came from which individuals and second by simply pooling all reads into a single sample, which substantially reduces run-time. Similarly, we ran Cufflinks in two ways: first running Cufflinks on each individual separately and using Cuffmerge Trapnell et al. (2012) to merge the reconstructed gene models into a single gene mode, and second by running Cufflinks on the pooled dataset. Given the gene model reconstructed by each of the four methods we used `altra` to estimate transcript expression values in each individual, at each locus.

We then tested each transcript for association with the splice eQTL SNP. For each transcript $t \in \{1, \ldots, T\}$ we performed a linear regression of the relative expression values $\boldsymbol{x}_t$ against the splice eQTL genotypes $\boldsymbol{g}$. Here the relative expression values $\boldsymbol{x}_t = (x_{t1}, \ldots, x_{tN})$ are given by $x_{ti} = \frac{\lambda_t^i l_t}{\sum_{t'=1}^{T} \lambda_{t'}^i l_{t'}}$ where $l_t$ is the transcript length. Each regression yields a $p$-value, $p_t$, and we used $p_{\min} = \min_t p_t$ as a test statistic for $H_0$ : SNP is unassociated with all $T$ transcripts at that locus. To assess significance of $p_{\min}$ we used permutation to get an empirical $p$-value (permuting the genotypes 1000 times for each locus, keeping the relative expressions fixed).

Overall, `altra` found more splice eQTLs than Cufflinks, both with and without pooling (Figure 8). For example, at a threshold of $p < 0.01$, with pooling `altra` found 30% of the

splice eQTLs, and Cufflinks found 12%. However, Cufflinks did identify some eQTLs missed by `altra`, as well as vice versa (Figure 7). The average performance of both methods was similar for pooled and non-pooled data; since `altra` is much faster for pooled data, pooling may be the preferred strategy in practice.

Neither `altra` nor Cufflinks identified the majority of the 187 splice eQTLs in Pickrell et al. (2010a). This may be partly explained by the fact that the original analysis Pickrell et al. (2010a) used a sensible and multi-step preprocessing procedure (averaging the fraction of reads mapping to each exon over technical replicates, normalizing, and correcting for confounding variables) that likely increased power. However, since some splice eQTLs are identified by `altra` and not by Cufflinks and vice versa, it appears that an exon-level analysis successfully identifies effects that are missed by current transcript-reconstruction methods.

# Conclusions

We presented `altra`, a novel Bayesian method for simultaneous transcript reconstruction and expression estimation using multiple RNA-Seq samples. A key novel feature of the method is that, given only a list of candidate splice sites (which can be obtained either from annotations or from the observed data), we specify a prior on transcripts that encourages different transcripts to share a similar structure, while at the same time allowing for common alternative splicing patterns, such as exon inclusion-exclusion, and alternative splice site usage. We found that the inferred transcripts from our method produced higher power to identify a set of known splice eQTLs than did inferred transcripts from Cufflinks, although neither method identified the majority of splice eQTLs.

Despite the innovations presented here, and more generally the very considerable research on transcript reconstruction Denoeud et al. (2008); Montgomery et al. (2010); Guttman et al. (2010); Filichkin et al. (2010); Trapnell et al. (2010); Zhao et al. (2011); Li et al. (2011a,b);

$-\log_{10}(p\text{-value})$

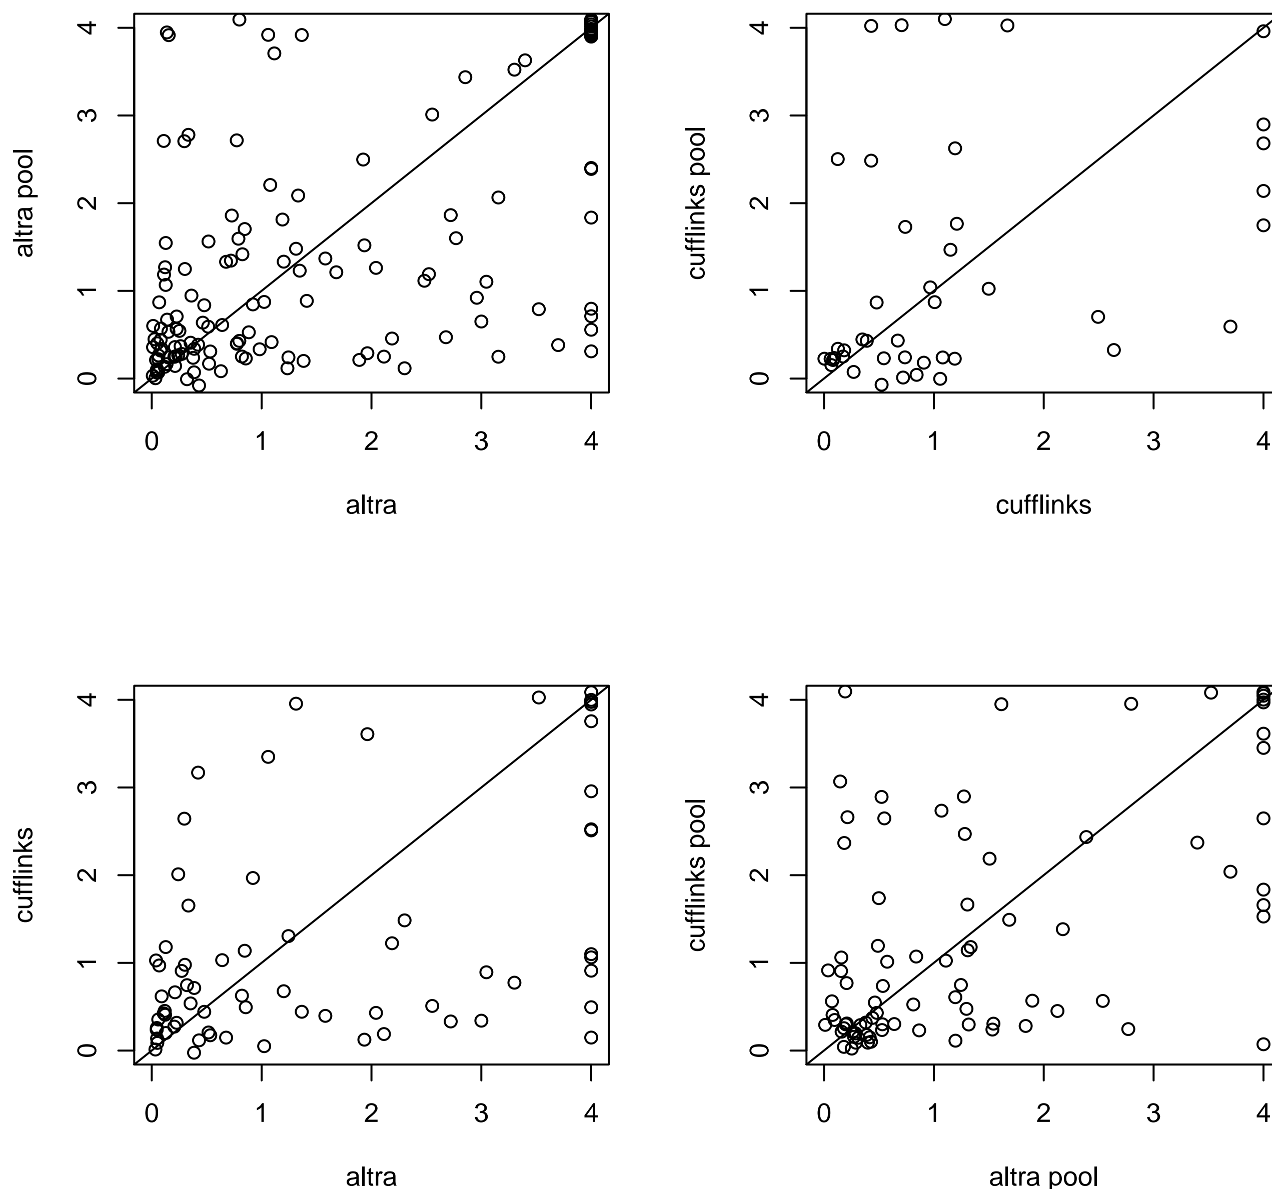


Figure 7: ***P*-values often differ appreciably across transcript reconstruction strategies at the same locus.** At each locus we ran: `altra` using data from all individuals simultaneously (altra); Cufflinks on each of the individuals and merged the resulting set of gene models with Cuffmerge (Cufflinks); `altra` on the pooled set of individuals (altra pool); Cufflinks on the pooled set of individuals (Cufflinks pool). We run `altra` with options -J 2 -M 2 -D 12 and $T = 5$ and Cufflinks (v2.1.1) with default options and reference annotation.

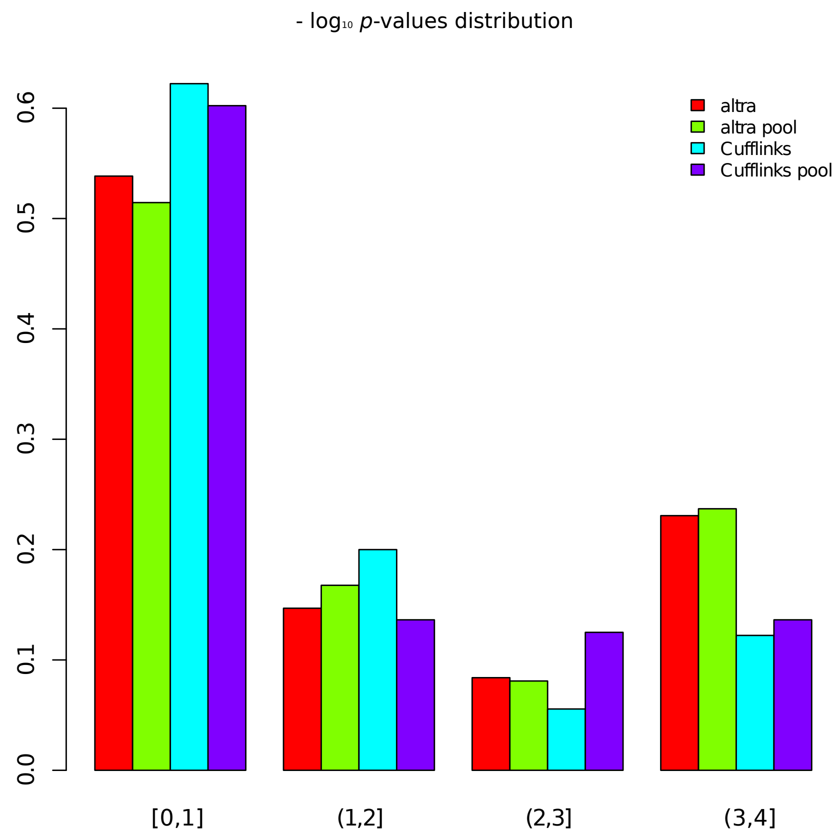


Figure 8: From left to right: fraction of empirical negative log 10 $p$-values from Figure 7 in intervals $[0, 1]$, $(1, 2]$, $(2, 3]$, $(3, 4]$.

Feng et al. (2011); Xia et al. (2011); Mezlini et al. (2012); Hiller and Wong (2013); Mangul et al. (2012); Lin et al. (2012); Behr et al. (2013); Bernard et al. (2013); Zeller et al. (2013), it is clear that transcript reconstruction remains a very challenging problem. Two aspects of our results particularly emphasize this. First, when we run our method on the same data multiple times we often infer quite different gene models that all appear broadly consistent with the observed data. Although finding different solutions in different runs may appear to be a defect of our method, and suggests lack of convergence of our stochastic inference procedure, there is no guarantee that deterministic methods that always produce the same answer for every run are more accurate - indeed it seems likely that, for the read lengths and coverage considered in the data here, the gene model is impossible to infer precisely from the data, and the convergence problems suffered by our method are, perhaps, a symptom of this (convergence is harder in problems with multiple plausible distinct solutions). Further, this may remain true even for larger datasets with longer reads in cases where the gene model is complex (many different alternative spliced patterns).

Second, when we attempted to map known splice eQTLs with estimated expression levels of inferred transcripts, we found that only a minority were identified by our analyses (using either Cufflinks or `altra`). This suggests that it would be unwise to rely exclusively on transcript-reconstruction-based methods in splice eQTL analyses, and that they should be complemented by event-based or exon-based analyses. (It is of course possible that transcript-based methods can identify splice eQTLs that are missed by exon-based analyses, but we did not assess this here).

Our methods are computationally intensive, particularly when used to jointly infer the gene model and transcript expression levels in many samples (e.g. 69 here). Computation is substantially reduced by first pooling the samples to infer the gene model, and then estimating expression levels as a second step. Although such two-stage procedures are in theory less attractive, in practice we found that pooling did not reduce accuracy (in terms of

identifying splice eQTLs). With this pooling strategy, our method took roughly 1.1 minutes to estimate the gene model, and 40 minutes to estimate transcript expression levels in each sample for a typical gene.

Although our models and methods can make use of gene annotation databases when available, these are not required. Our modeling approach requires only a list of putative 5' and 3' splice sites and a list of putative transcription start and end sites, all of which can be obtained directly from the RNA-Seq data: the putative splice sites can be obtained using a splice-aware mapper (e.g. Trapnell et al. (2009)) to map "junction" reads, and the putative start and end sites can be inferred from a segmentation method (e.g. Nowak et al. (2011)) that can detect sudden changes in the coverage pattern. Of these, the latter process is likely more error prone: for example, sudden changes in coverage could also be caused by mapping biases. Thus supplementing these kinds of data with lists of known transcription start and end sites could help improve accuracy.

Important features of RNA-Seq data not included in our model include sequence preference biases, positional biases (e.g. 3' bias), and mappability biases. Our model could be extended to include these biases, and ultimately this may be important for accurate gene model reconstruction and expression level estimation. For example, in *COMMD4* we saw that `altra` inferred low expression levels for the longest inferred transcripts, and higher expression for an unannotated transcript that excludes the first three exons, which could be due to unmodelled 3' bias in the sequence data. However, given the inherent difficulty of the problem, it is unclear whether the gains from incorporating these features would be sufficient to justify what could be considerable additional work.

Finally, we note that our modeling approach based on flexons opens a door to a different possible approach to the problem of understanding transcriptional processes. Specifically, our prior on the gene model introduces parameters associated with each flexon, to capture both the frequency of inclusion/exclusion of the flexon, and the frequency of use of alternative

3' and 5' splice sites. Here we used this model as a prior on transcripts, and used it to infer a small number of specific transcripts that are expressed in the data. However, an alternative would be to bypass transcript reconstruction altogether, and instead to estimate the parameters of the model ($\phi$) directly from the data, integrating over the full set of transcripts that could be generated by the prior, by replacing the Poisson mean in (1) with

$$x_r^i \mid \phi, \lambda_i, \epsilon \sim Pois\left((E[z_r(\tilde{T})|\phi]\lambda^i + \epsilon)C^i\right). \tag{5}$$

where $z_r(\tilde{T})$ is an indicator for whether a read of type $r$ is compatible with a transcript $\tilde{T}$, and the expectation $E[]$ is taken with respect to our "prior" distribution on transcripts $p(\tilde{T}|\phi)$. This model is motivated by an underlying assumption that each transcript produced at the locus is generated from the prior distribution $p(\tilde{T}|\phi)$ (rather than being generated by a discrete distribution on a small number of possible transcripts as we assume here). Questions of splicing differences among groups (e.g., in the case of splice eQTLs, differences among genotype groups) could be addressed by testing for differences in $\phi$ between groups. Computationally, this approach replaces the difficult combinatorial search over the space of possible gene models $\tau$, with the evaluation of an expectation that we believe will be more tractable. Statistically, this approach changes focus from reconstruction of individual transcripts to the average behavior across transcripts, particularly with regard to features of interest such as inclusion of a particular flexon or the use of a particular 3' splice site. This approach thus merges ideas from both transcript-reconstruction methods and "event-based" methods such as methods based on the "percentage spliced in" Katz et al. (2010), and we view it as a promising avenue for future development.

### List of abbreviations

eQTL: Expression quantitative trait loci; SNP: Single-nucleotide polymorphism.

### Competing interests

The authors declare that they have no competing interests.

### Author's contributions

MS conceived of the method, EP implemented and tested the method. All authors wrote, read and approved the final manuscript.

### Acknowledgements

We thank John Zekos for technical support and members of the Przeworski, Pritchard and Stephens labs for helpful discussions, particularly Stoyan Georgiev and Kevin Bullaughey. This work was supported by NIH grant HG02585 to MS.

### Additional Files

**Additional file 1**

This file contains a detailed description of some aspects of the method, specifically: read and transcript representation, *flexon* construction, implementation of the Metropolis Hastings algorithm.

## References

Anders, S. and Huber, W. (2010). "Differential expression analysis for sequence count data." *Genome Biology*, 11, 10, R106.

Anders, S., McCarthy, D. J., Chen, Y., Okoniewski, M., Smyth, G. K., Huber, W., and Robinson, M. D. (2013). "Count-based differential expression analysis of RNA sequencing data using R and Bioconductor." *Nat Protocols*, 8, 9, 1765–1786.

Anders, S., Reyes, A., and Huber, W. (2012). "Detecting differential usage of exons from RNA-Seq data." *Nature Precedings*.

Au, K. F., Jiang, H., Lin, L., Xing, Y., and Wong, W. H. (2010). "Detection of splice junctions from paired-end RNA-seq data by SpliceMap." *Nucleic Acids Res*, 38, 14, 4570–4578.

Barbosa-Morais, Nuno L. ad Irimia, M., Pan, Q., Xiong, H. Y., Gueroussov, S., Lee, L. J., Slobodeniuc, 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, 6114, 1587–1593.

Behr, J., Kahles, A., Zhong, Y., Sreedharan, V., Drewe, P., and Ratsch, G. (2013). "MI-TIE: Simultaneous RNA-Seq-based transcript identification and quantification in multiple samples." *BMC Bioinformatics*.

Bernard, E., Jacob, L., Mairal, J., and Vert, J. (2013). "Efficient RNA Isoform Identification and Quantification from RNA-Seq Data with Network Flows." Tech. rep., INRIA.

Blekhman, R., Marioni, J. C., Zumbo, P., Stephens, M., and Gilad, Y. (2010). "Sex-specific and lineage-specific alternative splicing in primates." *Genome Res*, 20, 2, 180–189.

Brooks, A. N., Yang, L., Duff, M. O., Hansen, K. D., Park, J. W., Dudoit, S., Brenner, S. E., and Graveley, B. R. (2011). "Conservation of an RNA regulatory map between Drosophila and mammals." *Genome Res*, 21, 2, 193–202.

Denoeud, F., Aury, J.-M., Da Silva, C., Noel, B., Rogier, O., Delledonne, M., Morgante, M., Valle, G., Wincker, P., Scarpelli, C., and et al. (2008). "Annotating genomes with massive-scale RNA sequencing." *Genome Biology*, 9, 12, R175.

Dillies, M.-A., Rau, A., Aubert, J., Hennequet-Antier, C., Jeanmougin, M., Servant, N., Keime, C., Marot, G., Castel, D., Estelle, J., and et al. (2012). "A comprehensive evaluation of normalization methods for Illumina high-throughput RNA sequencing data analysis." *Briefings in bioinformatics*.

Djebali, S., Davis, C. A., Merkel, A., Dobin, A., Lassmann, T., Mortazavi, A., Tanzer, A., Lagarde, J., Lin, W., Schlesinger, F., Xue, C., Marinov, G. K., Khatun, J., Williams, B. A., Zaleski, C., Rozowsky, J., Röder, M., Kokocinski, F., Abdelhamid, R. F., Alioto, T., Antoshechkin, I., Baer, M. T., Bar, N. S., Batut, P., Bell, K., Bell, I., Chakrabortty, S., Chen, X., Chrast, J., Curado, J., Derrien, T., Drenkow, J., Dumais, E., Dumais, J., Duttagupta, R., Falconnet, E., Fastuca, M., Fejes-Toth, K., Ferreira, P., Foissac, S., Fullwood, M. J., Gao, H., Gonzalez, D., Gordon, A., Gunawardena, H., Howald, C., Jha, S., Johnson, R., Kapranov, P., King, B., Kingswood, C., Luo, O. J., Park, E., Persaud, K., Preall, J. B., Ribeca, P., Risk, B., Robyr, D., Sammeth, M., Schaffer, L., See, L.-H. H., Shahab, A., Skancke, J., Suzuki, A. M. M., Takahashi, H., Tilgner, H., Trout, D., Walters, N., Wang, H., Wrobel, J., Yu, Y., Ruan, X., Hayashizaki, Y., Harrow, J., Gerstein, M., Hubbard, T., Reymond, A., Antonarakis, S. E., Hannon, G., Giddings, M. C., Ruan, Y., Wold, B., Carninci, P., Guigó, R., and Gingeras, T. R. (2012). "Landscape of transcription in human cells." *Nature*, 489, 7414, 101–108.

Ergun, A., Doran, G., Costello, J. C., Paik, H. H., Collins, J. J., Mathis, D., Benoist, C., and Consortium, I. (2013). "Differential splicing across immune system lineages." *Proc Natl Acad Sci*.

Feng, J., Li, W., and Jiang, T. (2011). "Inference of isoforms from short sequence reads." *J Comp Biol*, 18, 3, 305–21.

Filichkin, S. A., Priest, H. D., Givan, S. A., Shen, R., Bryant, D. W., Fox, S. E., Wong,

W.-K., and Mockler, T. C. (2010). "Genome-wide mapping of alternative splicing in Arabidopsis thaliana." *Genome Res*, 20, 1, 45–58.

Gonzàlez-Porta, M., Calvo, M., Sammeth, M., and Guigó, R. (2012). "Estimation of alternative splicing variability in human populations." *Genome Res*, 22, 3, 528–538.

Griebel, T., Zacher, B., Ribeca, P., Raineri, E., Lacroix, V., Guigó, R., and Sammeth, M. (2012). "Modelling and simulating generic RNA-Seq experiments with the flux simulator." *Nucleic Acids Res*, 40, 20, 10073–10083.

Guttman, M., Garber, M., Levin, J. Z., Donaghey, J., Robinson, J., Adiconis, X., Fan, L., Koziol, M. J., Gnirke, A., Nusbaum, C., and et al. (2010). "Ab initio reconstruction of cell type-specific transcriptomes in mouse reveals the conserved multi-exonic structure of lincRNAs." *Nature Biotechnol*, 28, 5, 503–510.

Hiller, D. and Wong, W. (2013). "Simultaneous Isoform Discovery and Quantification from RNA-Seq." *Stat Biosci*, 5:100–118.

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, 12, 1009–15.

Kurmangaliyev, Y. Z., Sutormin, R. A., Naumenko, S. A., Bazykin, G. A., and Gelfand, M. S. (2013). "Functional implications of splicing polymorphisms in the human genome." *Hum Mol Gen*, 22.

Lappalainen, T., Sammeth, M., Friedländer, M. R., 't Hoen, P. A. C., Monlong, J., Rivas, M. A., Gonzàlez-Porta, M., Kurbatova, N., Griebel, T., Ferreira, P. G., Barann, M., Wieland, T., Greger, L., van Iterson, M., Almlöf, J., Ribeca, P., Pulyakhina, I., Esser, D., Giger, T., Tikhonov, A., Sultan, M., Bertier, G., MacArthur, D. G., Lek, M., Lizano, E.,

Buermans, H. P. J., Padioleau, I., Schwarzmayr, T., Karlberg, O., Ongen, H., Kilpinen, H., Beltran, S., Gut, M., Kahlem, K., Amstislavskiy, V., Stegle, O., Pirinen, M., Montgomery, S. B., Donnelly, P., McCarthy, M. I., Flicek, P., Strom, T. M., Consortium, T. G., Lehrach, H., Schreiber, S., Sudbrak, R., Carracedo, A., Antonarakis, S. E., Häsler, R., Syvänen, A.-C., van Ommen, G.-J., Brazma, A., Meitinger, T., Rosenstiel, P., Guigó, R., Gut, I. G., Estivill, X., and Dermitzakis, E. T. (2013). "Transcriptome and genome sequencing uncovers functional variation in humans." *Nature*, 501, 7468, 506–511.

Li, J. J., Jiang, C.-R., Brown, J. B., Huang, H., and Bickel, P. J. (2011a). "Sparse linear modeling of next-generation mRNA sequencing (RNA-Seq) data for isoform discovery and abundance estimation." *Proc Natl Acad Sci*, 108, 50, 19867–72.

Li, W., Feng, J., and Jiang, T. (2011b). "IsoLasso: A LASSO Regression Approach to RNA-Seq Based Transcriptome Assembly." *Journal of computational biology a journal of computational molecular cell biology*, 18, 11, 1693–707.

Lin, Y.-Y., Dao, P., Hach, F., Bakhshi, M., Mo, F., Lapuk, A., Collins, C., and Sahinalp, S. (2012). "CLIIQ: Accurate comparative detection and quantification of expressed isoforms in a population." In *Proceedings of the 12th international conference on Algorithms in Bioinformatics*, vol. 7534 of *Lecture Notes in Computer Science*, 178–189.

Mangul, S., Caciula, A., Glebova, O., Mandoiu, I., and Zelikovsky, A. (2012). "Improved transcriptome quantification and reconstruction from RNA-Seq reads using partial annotations." *In Silico Biology*, 11, 5-6.

Merkin, J., Russell, C., Chen, P., and Burge, C. B. (2012). "Evolutionary dynamics of gene and isoform regulation in mammalian tissues." *Science*, 338, 6114, 1593–1599.

Mezlini, A. M., Smith, E. J. M., Fiume, M., Buske, O., Savich, G., Shah, S., Aparicion,

S., Chiang, D., Goldenberg, A., and Brudno, M. (2012). “iReckon: Simultaneous isoform discovery and abundance estimation from RNA-seq data.” *Genome Res*.

Montgomery, S. B., Sammeth, M., Gutierrez-Arcelus, M., Lach, R. P., Ingle, C., Nisbett, J., Guigo, R., and Dermitzakis, E. T. (2010). “Transcriptome genetics using second generation sequencing in a Caucasian population.” *Nature*, 464, 7289, 773–777.

Mortazavi, A., Williams, B. A., McCue, K., Schaeffer, L., and Wold, B. (2008). “Mapping and quantifying mammalian transcriptomes by RNA-Seq.” *Nature Methods*, 5, 7, 621–628.

Nowak, G., Hastie, T., Pollack, J. R., and Tibshirani, R. (2011). “A fused lasso latent feature model for analyzing multi-sample aCGH data.” *Biostatistics Oxford England*, 12, 4, 776–791.

Pan, Q., Shai, O., Lee, L. J., Frey, B. J., and Blencowe, B. J. (2008). “Deep surveying of alternative splicing complexity in the human transcriptome by high-throughput sequencing.” *Nat Genet*, 40, 12, 1413–1415.

Pickrell, J. K., Marioni, J. C., Pai, A. A., Degner, J. F., Engelhardt, B. E., Nkadori, E., Veyrieras, J.-B., Stephens, M., Gilad, Y., and Pritchard, J. K. (2010a). “Understanding mechanisms underlying human gene expression variation with RNA sequencing.” *Nature*, 464, 7289, 768–72.

Pickrell, J. K., Pai, A. A., Gilad, Y., and Pritchard, J. K. (2010b). “Noisy splicing drives mRNA isoform diversity in human cells.” *PLoS Genetics*, 6, 12, 8.

Reyes, A., Anders, S., Weatheritt, R. J., Gibson, T. J., Steinmetz, L. M., and Huber, W. (2013). “Drift and conservation of differential exon usage across tissues in primate species.” *Proc Natl Acad Sci*, 110, 38, 15377–15382.

Roberts, A., Pimentel, H., Trapnell, C., and Pachter, L. (2011). "Identification of novel transcripts in annotated genomes using RNA-Seq." *Bioinformatics*.

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, 1, 139–40.

Salzman, J., Jiang, H., and Wong, W. H. (2010). "Statistical Modeling of RNA-Seq Data." Tech. rep., Stanford University.

Seok, J., Xu, W., Gao, H., Davis, R. W., and Xiao, W. (2012). "JETTA: Junction and Exon Toolkits for Transcriptome Analysis." *Bioinformatics*, 28, 9, 1274–1275.

Sun, W. and Hu, Y. (2013). "eQTL Mapping Using RNA-seq Data." 5, 1, 198–219.

Trapnell, C., Pachter, L., and Salzberg, S. L. (2009). "TopHat: discovering splice junctions with RNA-Seq." *Bioinformatics*, 25, 9, 1105–1111.

Trapnell, C., Roberts, A., Goff, L., Pertea, G., Kim, D., Kelley, D. R., Pimentel, H., Salzberg, S. L., Rinn, J. L., and Pachter, L. (2012). "Differential gene and transcript expression analysis of RNA-seq experiments with TopHat and Cufflinks." *Nat. Protocols*, 7, 3, 562–578.

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 Biotechnol*, 28, 5, 511–515.

Turro, E., Su, S.-Y., Goncalves, A., Coin, L. J., Richardson, S., and Lewin, A. (2011). "Haplotype and isoform specific expression estimation using multi-mapping RNA-seq reads." *Genome Biology*, 12, 2, R13.

Wang, E. T., Sandberg, R., Luo, S., Khrebtukova, 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.

Xia, Z., Wen, J., Chang, C.-C., and Zhou, X. (2011). "NSMAP: A method for spliced isoforms identification and quantification from RNA-Seq." *BMC Bioinformatics*, 12, 1, 162.

Zeller, G., Goernitz, N., Kahles, A., Behr, J., Mudrakarta, P., Sonnenburg, S., and Raetsch, G. (2013). "mTim: Rapid and accurate transcript reconstruction from RNA-Seq data." *ArXiv e-prints*.

Zhao, K., Lu, Z. X., Park, J. W., Zhou, Q., and Xing, Y. (2013). "GLiMMPS: Robust statistical model for regulatory variation of alternative splicing using RNA-Seq data." *Genome Biology*, 14, 7, R74+.

Zhao, Z., Chi, T. N., Deng, N., Johnson, K. M., and Zhu, D. (2011). "SPATA: a seeding and patching algorithm for de novo transcriptome assembly." *IEEE International Conference on Bioinformatics and Biomedicine Workshop (IEEE BIBMW'11)*, 26–33.