

UCLA

UCLA Previously Published Works

Title

Applying meta-analysis to genotype-tissue expression data from multiple tissues to identify eQTLs and increase the number of eGenes

Permalink

<https://escholarship.org/uc/item/467476mx>

Journal

Bioinformatics, 33(14)

ISSN

1367-4803

Authors

Duong, Dat
Gai, Lisa
Snir, Sagi
et al.

Publication Date

2017-07-15

DOI

10.1093/bioinformatics/btx227

Peer reviewed

Applying meta-analysis to genotype-tissue expression data from multiple tissues to identify eQTLs and increase the number of eGenes

Dat Duong,^{1,*} Lisa Gai,¹ Sagi Snir,^{2,3} Eun Yong Kang,¹ Buhm Han,^{4,5}
Jae Hoon Sul^{6,†} and Eleazar Eskin^{1,7,*†}

¹Department of Computer Science, University of California, Los Angeles, CA 90095, USA, ²Institute of Evolution, University of Haifa, Haifa 3498838, Israel, ³Department of Evolutionary and Environmental Biology, University of Haifa, Haifa 3498838, Israel, ⁴Department of Convergence Medicine, University of Ulsan College of Medicine, Seoul 05505, Republic of Korea, ⁵Asan Institute for Life Sciences, Asan Medical Center, Seoul 05505, Republic of Korea, ⁶Department of Psychiatry and Biobehavioral Sciences, University of California, Los Angeles, CA 90095, USA and ⁷Department of Human Genetics, University of California, Los Angeles, CA 90095, USA

*To whom correspondence should be addressed.

†The authors wish it to be known that, in their opinion, the last two authors should be regarded as Joint Authors.

Abstract

Motivation: There is recent interest in using gene expression data to contextualize findings from traditional genome-wide association studies (GWAS). Conditioned on a tissue, expression quantitative trait loci (eQTLs) are genetic variants associated with gene expression, and eGenes are genes whose expression levels are associated with genetic variants. eQTLs and eGenes provide great supporting evidence for GWAS hits and important insights into the regulatory pathways involved in many diseases. When a significant variant or a candidate gene identified by GWAS is also an eQTL or eGene, there is strong evidence to further study this variant or gene. Multi-tissue gene expression datasets like the Gene Tissue Expression (GTEx) data are used to find eQTLs and eGenes. Unfortunately, these datasets often have small sample sizes in some tissues. For this reason, there have been many meta-analysis methods designed to combine gene expression data across many tissues to increase power for finding eQTLs and eGenes. However, these existing techniques are not scalable to datasets containing many tissues, like the GTEx data. Furthermore, these methods ignore a biological insight that the same variant may be associated with the same gene across similar tissues.

Results: We introduce a meta-analysis model that addresses these problems in existing methods. We focus on the problem of finding eGenes in gene expression data from many tissues, and show that our model is better than other types of meta-analyses.

Availability and Implementation: Source code is at <https://github.com/datduong/RECOV>.

Contact: eeskin@cs.ucla.edu or datdb@cs.ucla.edu

Supplementary information: Supplementary data are available at *Bioinformatics* online.

1 Introduction

Expression quantitative trait loci (eQTLs) studies find eQTLs, which are genetic variants associated with gene expression, and eGenes, which are genes whose expression levels are associated with at least one genetic variant. eQTL studies are related to traditional genome-wide association studies (GWAS) which find variants associated with disease.

Both eQTL studies and GWAS focus on single nucleotide polymorphisms (SNPs). Many SNPs found by GWAS are located in intergenic regions, and their relationship to the disease phenotype is often not obvious. Gene expression is an intermediate phenotype between a causal SNP and a disease (Huang *et al.*, 2014). Thus, eQTL studies may provide biological insights into the mechanism through which disease occurs. If a significant SNP identified by GWAS is found to be an eQTL, there is a strong evidence to further study the

variant. For this reason, top hits in GWAS that are also eQTLs are of special interest. In fact, recent GWAS have confirmed that many disease-causing variants are eQTLs (Albert, 2016; Liu et al., 2016; Nieuwenhuis et al., 2016). Similarly, genes near GWAS-significant SNPs that are identified as eGenes may warrant further study as candidate causal genes. Thus, eQTL studies provide great supporting evidence for GWAS results and important insights into the regulatory pathways involved in many diseases.

The underlying approach behind eQTL studies and GWAS is the same. In an eQTL study, one performs association tests between the genotype data and the gene expression (instead of disease statuses) to identify variants that are associated with the gene expression. eQTLs and eGenes may be specific to only one or a group of tissues, as a gene is not always uniformly expressed in every tissue. For example, SNPs associated with schizophrenia have been found to be eQTLs in only the brain tissues, indicating that schizophrenia affects how the brain functions (Fromer et al., 2016). For this reason, there have been recent large-scale studies to collect gene expression data in many tissues, such as the Genotype-Tissue Expression (GTEx) project (The GTEx Consortium, 2015). This GTEx dataset contains gene expression data in 44 tissues and genotypes of 5 million SNPs for over 300 individuals.

To find eQTLs from the GTEx data and other multi-tissue datasets, one can apply the traditional tissue-by-tissue (TBT) approach, in which a separate eQTL study is done for each tissue. However, many tissues do not have enough samples to detect SNPs that are weakly associated with the gene expressions. To address this issue, there have been many efforts in developing different types of meta-analysis, which gather data from many tissues to increase the total sample size and power to detect eQTLs. Two notable methods are Meta-Tissue and eQTLBma. Both have been shown to outperform the traditional TBT method (Flutre et al., 2013; Sul et al., 2013).

Meta-Tissue and eQTLBma have an important limitation that reduces their applicability to large gene expression datasets such as the GTEx data. Both methods are computationally intensive and should be used for datasets containing at most 10 or 20 tissues, respectively (Flutre et al., 2013; Sul et al., 2013). Meta-Tissue uses both linear mixed models (LMMs) and fixed (or random) effects meta-analysis to combine data from many tissues. Meta-Tissue must estimate the variance components in its LMM setup for every pair of variant and gene expression; thus, its runtime is impractical when there are thousands of genes or too many tissues (Sul et al., 2013).

eQTLBma uses a Bayesian approach that considers all possible combinations of tissues in which a SNP is an eQTL. This setup corresponds to 2^T configurations where T is the number of tissues, making the method infeasible when T is 44 like in the GTEx data (Flutre et al., 2013).

As an alternative to Meta-Tissue and eQTLBma, the GTEx consortium used a meta-analysis software called Metasoft, introduced by Han and Eskin (2011). Metasoft is equivalent to Meta-Tissue without the LMM setup (Han and Eskin, 2011, 2012). Metasoft extends the random effects (REs) meta-analysis model; this extended model is named RE2 (Han and Eskin, 2011).

However, Meta-Tissue, eQTLBma and RE2 assume that a SNP has independent effect on a given gene's expression in each tissue. This ignores the fact that the same SNP tends to have similar effects in related tissues (The GTEx Consortium, 2015).

Recently, Acharya et al. (2016) introduced a method that amends this shortcoming in Meta-Tissue, eQTLBma and RE2. The model developed by Acharya et al. (2016) requires genotype and gene expression data for each individual in each tissue. Their implementation in R, using the JARGUAR library, requires loading all these data into memory. When there are many genes and tissues, this approach can be memory intensive.

In this article, we present a novel meta-analysis method named RECOV. Unlike Meta-Tissue and eQTLBma, RECOV is applicable to large gene expression datasets and can analyze all 44 tissues in the GTEx data. Like JARGUAR, RECOV considers the biological insight that a variant may have similar effects on a gene across tissues. However, unlike JARGUAR, RECOV needs only the summary statistic (i.e. SNP effect and its variance) at each SNP in each tissue and not the complete genotype and gene expression data for each individual. RECOV is based on the RE2 meta-analysis framework and uses a covariance (COV) matrix to explicitly model the correlation of a SNP effect on the same gene's expression in similar tissues.

In the Section 2, we describe RECOV in detail and demonstrate how it can be used to identify eGenes from eQTL studies in more than one tissue. In the Results section, we use simulated datasets to show that RECOV has correct false positive rate (FPR). We then apply RECOV to real multi-tissue expression data from the GTEx dataset. Our results show that RECOV detects more eGenes than previous RE2 and TBT methods.

2 Materials and methods

We begin by introducing the notations in this article. We use $\mathbf{x} \in R^n$ to specify a vector \mathbf{x} with dimension n , and $\mathbf{Z} \in R^{n \times m}$ to specify a matrix \mathbf{Z} with dimension $n \times m$. We use x_i to denote the i th element in \mathbf{x} , and likewise, Z_{ij} to specify entry ij in \mathbf{Z} . We denote an item k in the set K by $k \in K$, and a set $\{a_1 \dots a_K\}$ indexed by k by using $\{a_k\}_{k \in K}$, where the subscript $k \in K$ is omitted whenever the context is clear. The size of the set K is denoted as $|K|$.

2.1 Detecting one eGene via an eQTL study

2.1.1 eQTL study in one tissue

We begin with an eQTL study in one tissue t . An eQTL study finds every eQTL associated with the expression level of a specific gene g . To do this, the study tests each variant v in the set V against the expression of g in a sequential fashion. To set up the problem, suppose we represent the gene expression for m individuals in tissue t as a vector $\mathbf{q} \in R^m$, and we want to find the effect of variant v on g . Let $\mathbf{s} \in R^m$ be the standardized genotypes of this v . The eQTL study assumes the following model

$$\mathbf{q} = \beta_{vgt}\mathbf{s} + \epsilon_{vgt} \quad (1)$$

where $\epsilon \in R^m$ is the vector of sampling errors $\epsilon \sim N(0, \sigma_\epsilon^2 \mathbf{I})$, and $\beta_{vgt} \in R$ is the true effect size of the variant v on g in tissue t (Eskin, 2015). The estimate b_{vgt} of the true value β_{vgt} can be computed using the basic least squares equation $b_{vgt} = \argmin_{\beta_{vgt}} \|\mathbf{q} - \beta_{vgt}\mathbf{s}\|_2^2$ (Abraham and Ledolter, 2006). This solution is

$$b_{vgt} = (\mathbf{s}^\top \mathbf{s})^{-1} \mathbf{s}^\top \mathbf{q} \quad \text{where} \quad b_{vgt} \sim N\left(\beta_{vgt}, (\mathbf{s}^\top \mathbf{s})^{-1} \sigma_\epsilon^2\right) \quad (2)$$

By using Equation (2) and writing the null hypothesis $H_0: \beta_{vgt} = 0$, one can do a hypothesis test to assert if v has an effect on g . To do this test, compute the estimate $\hat{\sigma}_{\epsilon_{vgt}}^2$ of $\sigma_{\epsilon_{vgt}}^2$ by

$$\hat{\sigma}_{\epsilon_{vgt}}^2 = \frac{1}{m-1} \sum_{i=1}^m (q_i - b_{vgt}s_i)^2 \quad (3)$$

and estimate the variance d_{vgt} of b_{vgt} by

$$d_{vgt} = (\mathbf{s}^\top \mathbf{s})^{-1} \hat{\sigma}_{\epsilon_{vgt}}^2 \quad (4)$$

then compute the p -value $p_{vgt} = p - \text{value}(b_{vgt})$ (Abraham and Ledolter, 2006; Eskin, 2015). If p_{vgt} is less than some significance level, then we reject $H_0: \beta_{vgt} = 0$, and conclude that v is an eQTL of

g in tissue t . When many variants are tested, we must apply a multiple testing correction; for example, one can apply Bonferroni correction by using the threshold $\alpha/|V|$ where α is the significance level for the whole family of tests. The Bonferroni correction is conservative when there is linkage disequilibrium (LD) in set V . There exist other methods that can handle LD better than the Bonferroni correction (Darnell *et al.*, 2012; Joo *et al.*, 2014, 2016; Hormozdiari *et al.*, 2015).

2.1.2 Using an eQTL study in one tissue to discover one eGene

Because an eQTL study tests each variant $v \in V$ against gene g in a tissue t , from a single eQTL study we have a set of p -values $\{p_{vgt}\}_{v \in V}$. The minimum $p_{gt} = \min_{v \in V} \{p_{vgt}\}$ is defined to be the observed eGene statistic at gene g in tissue t (The GTEx Consortium, 2015). Define $\alpha_{p_{gt}} = p - \text{value}(p_{gt})$ to be the eGene p -value (The GTEx Consortium, 2015). The eGene p -value depends on two important factors: the number of variants $|V|$, and the LD of the variants. In practice, $\alpha_{p_{gt}}$ is computed by doing a permutation test (Sul *et al.*, 2015; Duong *et al.*, 2016). In brief, in the k th permutation, one would permute the gene expressions among the individuals, and compute a new $p_{gt}^{(k)} = \min_{v \in V} \{p_{vgt}^{(k)}\}$. $\alpha_{p_{gt}}$ is the ranking of the observed p_{gt} with respect to the density created from many $p_{gt}^{(k)}$. If its eGene p -value $\alpha_{p_{gt}}$ is less than some desired threshold, one can then conclude that g is an eGene in tissue t .

2.2 TBT analysis to find one or many eGenes

When there are genotype-tissue expression data from many tissues, TBT analysis is the standard method to find eGenes (Sul *et al.*, 2013; The GTEx Consortium, 2015). TBT tests whether or not the gene has at least one eQTL in each tissue by examining each tissue individually. Suppose the gene is expressed in T tissues. Then TBT performs T eQTL studies (one test in each tissue). The null hypothesis is that the gene is not an eGene in any tissue. This hypothesis is equivalent to saying that no eQTL is found for this gene in any tissue.

Three layers of multiple testing correction are required since TBT performs one test per gene in each tissue. The first layer of multiple testing correction is applied within a tissue and corrects for LD of the variants tested against the gene. This correction can be done by using the permutation test to compute the eGene p -value for the gene in the tissue (Duong *et al.*, 2016; Sul *et al.*, 2013, 2015).

The second layer of multiple testing correction adjusts for the fact that we may test more than one gene within a tissue. For example, the GTEx pilot study tested thousands of genes within one tissue, and then transformed eGene p -values into eGene q -values to control for this multiple testing (Dabney *et al.*, 2010; The GTEx Consortium, 2015). This second layer of multiple testing correction is not needed if only one gene is tested in each tissue.

The third layer of multiple testing correction takes into account the fact that one gene is tested T number of times (once per tissue) (Sul *et al.*, 2013). In this article, we apply Bonferroni correction so that the false-positive threshold for any eGene q -value in each tissue is α/T , where α is 5% for example. In this layer, other multiple testing correction methods such as the Benjamini-Hochberg correction can be used instead of the Bonferroni correction. However, this paper focuses on the meta-analysis model, and measuring the performances of various multiple testing correction methods is beyond its scope.

2.3 Meta-analysis models for combining eQTL studies across tissues

We motivate the application of meta-analysis for combining eQTL studies across tissues. An eQTL is defined not only with respect to a gene, but also with respect to the tissue in which the gene expression is measured. eQTL studies of the same gene have been analyzed separately at the tissues level (The GTEx Consortium, 2015). We can better detect the effect of a variant on the gene by combining eQTL results across many tissues and modeling the relatedness of the effect sizes of one variant among the tissues.

When using meta-analysis to find many eGenes, it is important to emphasize that one would need only two layers of multiple testing correction. The first layer is applied within a gene to correct for LD because one tests many variants against the gene. The second layer is applied at the gene level because there is usually more than one gene being tested.

We define the notations to be used later. Suppose we have T eQTL studies (one study per tissue) that test the association of a variant v at a gene g . As before, denote the effect of this variant in the study (i.e. tissue) t as b_{vgt} , where b_{vgt} is computed using Equation (2). Denote the variance of b_{vgt} in the study t as d_{vgt} where d_{vgt} is computed using Equation (4). Let $\mathbf{b}_{vg} \in R^T$ contain the effects in these T studies, so that $\mathbf{b}_{vg} = [b_{vg1} \ b_{vg2} \ \dots \ b_{vgT}]^T$. Let $\mathbf{D}_{vg} = \text{diag}(d_{vg1} \dots d_{vgT})$.

2.3.1 REs and the RE2 model

The maximum likelihood procedure in the RE model assumes that \mathbf{b}_{vg} has the form (Han and Eskin, 2011; Thompson and Sharp, 1997)

$$\mathbf{b}_{vg} = \lambda_{vg} + \epsilon_{vg} \quad (5)$$

The random sampling errors ϵ_{vg} are estimated from the data and assumed to be $\epsilon_{vg} \sim N(0, \mathbf{D}_{vg})$. $\lambda_{vg} \in R^T$ in the RE model is a random variable, i.e. $\lambda_{vg} \sim N(\mu_{vg} \mathbf{1}, \tau_{vg}^2 \mathbf{I})$ with $\mu_{vg} \in R$ and $\tau_{vg}^2 \in R_+$. Here the number 1 denotes a vector with all entries equal to 1. The effect λ_{vg} is thus known as the random effect. μ_{vg} is the common true underlying effect that all the studies inherit. The term τ_{vg}^2 is the heterogeneity among the effects of the variant in T tissues.

Clearly, \mathbf{b}_{vg} comes from the distribution

$$\mathbf{b}_{vg} \sim N(\mu_{vg} \mathbf{1}, \tau_{vg}^2 \mathbf{I} + \mathbf{D}_{vg}) \quad (6)$$

The traditional RE model assumes that if the effect of the variant does not exist in any tissue, then $\mu_{vg} = 0$. However, it has been shown that this traditional null hypothesis does not yield optimal statistical power in detecting eQTLs (Han and Eskin, 2011, 2012). For this reason, the RE2 model assumes a different null hypothesis, that if the effect of the variant does not exist in any tissue, then $\mu_{vg} = 0$ and $\tau_{vg}^2 = 0$. The fact that $\tau_{vg}^2 = 0$ is a result of $\mu_{vg} = 0$ because when the effect does not exist, its variance must not exist (Han and Eskin, 2011, 2012; Kang *et al.*, 2014). We will compare our method against the RE2 model.

The null hypothesis H_0 in RE2 is

$$H_0 : \quad \mu_{vg} = 0 \quad \tau_{vg}^2 = 0 \quad (7)$$

The log likelihood ratio for testing this hypothesis becomes

$$\text{llr}_{vg} = 2 \log \frac{\sup_{\mu_{vg}, \tau_{vg}^2} L(\mathbf{b}_{vg} | \mu_{vg}, \tau_{vg}^2)}{L(\mathbf{b}_{vg} | \mu_{vg} = 0, \tau_{vg}^2 = 0)} \quad (8)$$

The function L denotes the likelihood function of the random variable \mathbf{b}_{vg} . The numerator $\sup_{\mu_{vg}, \tau_{vg}^2} L(\mathbf{b}_{vg} | \mu_{vg}, \tau_{vg}^2)$ may be estimated

using numerical methods or other heuristic methods. Here, we apply the Nelder-Mead method, which is a heuristic derivative-free search method.

In finding the supremum, one implicitly enforces $\tau_{vg}^2 \geq 0$. Due to this restricted parameter space, the asymptotic density of the likelihood ratio is an average of a χ_1^2 and χ_2^2 (Self and Liang, 1987; Han and Eskin, 2011). To find the p -value of this likelihood ratio when T is large, one can use this asymptotic density.

Otherwise, one may compute the likelihood ratio p -value by creating a density of likelihood ratios under the null hypothesis and ranking the observed likelihood ratio with respect to this density. This null density is made by sampling many instances of \mathbf{b}_{vg} using Equation (6) with $\mu_{vg} = \tau_{vg}^2 = 0$, and computing their corresponding llr_{vg} . If the p -value of llr_{vg} is significant, then v is an eQTL with respect to g in at least one tissue. Because we have 44 tissues in the GTEx data, we will use the asymptotic distribution of the likelihood ratio.

2.3.2 RECOV: REs model with a COV term

Here we present an extension to the RE model. We first discuss the COV term. Equation (5) of the RE model assumes $\lambda_{vg} \sim N(\mu_{vg} \mathbf{1}, \tau_{vg}^2 \mathbf{I})$ so that the effects of variant v toward gene g are independent across the tissues. However, tissues from the same body part are similar; in fact, many eQTLs are found to be consistent among many tissues (Flutre et al., 2013). From this observation, we must acknowledge that $\lambda_{vg} \sim N(\mu_{vg} \mathbf{1}, \Sigma_{vg})$ where Σ_{vg} is not diagonal. The term $\Sigma_{vg} \in R^{T \times T}$ models the COV of effect sizes of v among tissues conditioned on the gene g . The matrix Σ_{vg} contains $T \times T$ unknown parameters which are to be estimated. In practice, one has to assume a simpler form for Σ_{vg} . Here, we assume $\Sigma_{vg} \approx c_{vg} \mathbf{U}_{vg}$. The matrix \mathbf{U}_{vg} is estimated from the data. The term $c_{vg} \geq 0$ is an unknown scaling constant and is optimized jointly with the mean of regression coefficient μ_{vg} .

In this article, we compute the \mathbf{U}_{vg} at each variant-gene pair as follows. Denote $\mathbf{B}_g = [\mathbf{b}_{1g} \ \mathbf{b}_{2g} \ \cdots \ \mathbf{b}_{|V|g}]$ so that $\mathbf{B}_g \in R^{T \times |V|}$. A column in \mathbf{B}_g contains the effects of a variant in 44 tissues. To avoid reusing the data when testing a single SNP, we remove its effects in the 44 tissues when estimating its COV term. To do this, we divide all cis-variants of g into 10 separate segments according to their physical locations on the chromosome, and use the 9 segments that do not contain v to compute \mathbf{U}_{vg} . In particular, denote \mathbf{B}_{-vg} as the matrix \mathbf{B}_g without the effect sizes of the variants that belong to the same segment as v . \mathbf{U}_{vg} can be estimated as $\mathbf{U}_{vg} = \mathbf{B}_{-vg} \mathbf{B}_{-vg}^\top$ (after proper scaling is applied to \mathbf{B}_{-vg}). This computation is similar to how one would compute a kinship matrix using the genotype matrix (Eskin, 2015). In this scheme, we observe that the variants in strong LD with v are also removed, so that there are fewer vectors in \mathbf{B}_{-vg} that resemble \mathbf{b}_{vg} when computing \mathbf{U}_{vg} . This further helps reducing the problem of data reusing. Supplementary Table S2 shows that by removing SNPs in the same segment as v , we retain fewer SNPs that are in strong LD with v .

Now, we are ready to introduce this COV term \mathbf{U}_{vg} to the RE model. We extend the RE model so that when testing a variant v against gene g , we have

$$\mathbf{b}_{vg} = \lambda_{vg} + \epsilon_{vg} \quad (9)$$

where

$$\lambda_{vg} \sim N(\mu_{vg} \mathbf{1}, c_{vg} \mathbf{U}_{vg}) \quad \epsilon_{vg} \sim N(0, \mathbf{D}_{vg}) \quad (10)$$

Like before, the matrix \mathbf{D}_{vg} is known because it contains the observed variances of the SNP effects estimated by Equation (4).

This form for \mathbf{D}_{vg} is standard in meta-analysis (Thompson and Sharp, 1997). We now have

$$\text{cov}(\mathbf{b}_{vg}) = c_{vg} \mathbf{U}_{vg} + \mathbf{D}_{vg} \quad (11)$$

$$\mathbf{b}_{vg} \sim N(\mu_{vg} \mathbf{1}, \text{cov}(\mathbf{b}_{vg})) \quad (12)$$

The null hypothesis that v does not affect g in all T tissues is

$$\mathbf{H}_0 : \quad \mu_{vg} = 0 \quad c_{vg} = 0 \quad (13)$$

The alternative hypothesis implies that v has an effect in at least one of the T tissues.

Under this setting, the log likelihood ratio to test the hypothesis becomes

$$\text{llr}_{vg} = 2 \log \frac{\sup_{\mu_{vg}, c_{vg}} L(\mathbf{b}_{vg} | \mu_{vg}, c_{vg})}{L(\mathbf{b}_{vg} | \mu_{vg} = 0, c_{vg} = 0)} \quad (14)$$

Like in the RE model, in finding the supremum in the alternative, one enforces $c_{vg} \geq 0$. Due to this restricted parameter space, the asymptotic density of the likelihood ratio is an average of a χ_1^2 and χ_2^2 . Alternatively, one can compute the empirical p -value of this likelihood ratio with a permutation test. In any case, if the p -value of the likelihood ratio is significant, then v is an eQTL with respect to g in at least one tissue.

2.4 Using meta-analysis of eQTL in many tissues to identify eGenes

In practice, a set of variants V is tested against g . Here we describe how one can combine the meta-analysis result at each variant $v \in V$ to determine if g is an eGene.

Define $p_{vg} = p\text{-value}(\text{llr}_{vg})$ so that from many variants, we have the set of p -values $\{p_{vg}\}_{v \in V}$. The observed statistic at gene g is $p_g = \min_{v \in V} \{p_{vg}\}$. To determine if p_g is significant, one needs to compute its eGene p -value α_{p_g} (The GTEx Consortium, 2015). To control for multiple testing when LD exists between the variants, one can compute α_{p_g} using a permutation test (Duong et al., 2016; Sul et al., 2015; The GTEx Consortium, 2015). The permutation test creates a distribution of the observed p_g under the null hypothesis, which can then be used to compute the eGene p -value α_{p_g} of p_g .

This permutation test can be done as follows. Let K be the number of permutations. In the k th permutation, permute the gene expression of g among the individuals in each of the T tissues so that there are T permuted datasets. This permutation reflects the hypothesis that the gene is not an eGene in any tissue. Next, redo the meta-analysis at each variant $v \in V$ so that a new $p_g^{(k)} = \min_{v \in V} \{p_{vg}^{(k)}\}$ is computed. α_{p_g} is the fraction of times the observed p_g is less than $p_g^{(k)}$. The gene g is an eGene in at least one tissue if its eGene p -value α_{p_g} is below some threshold α .

In the pilot GTEx analysis, a set of genes G is being tested at once, so that one has a set $\alpha_G = \{\alpha_{p_g}\}_{g \in G}$. To control for the family wise error rate, one can apply Bonferroni correction to get the threshold $\alpha/|G|$. Any gene $g \in G$ with $\alpha_{p_g} < \alpha/|G|$ is an eGene in at least one of the T tissues.

2.4.1 Estimating eGene p -value

The permutation test above must be performed at every pair of variant $v \in V$ and gene $g \in G$ in a tissue t . The entire permutation test requires $K|V||G|T$ permuted datasets, which is time consuming. Here, we introduce an alternate method to estimate the eGene p -value. In essence, the permutation test estimates a function f that

maps a test statistic to its p -value. There is evidence that the correlation of test statistics at two variants is equal to their LD (Han et al., 2009). This holds true when the test statistics are effect sizes (Han et al., 2009).

In our meta-analysis, to properly estimate the eGene p -value α_{p_g} , we must consider the effect of LD in the set of variants V on the observed statistic at each variant v . At each variant, it does not matter whether we treat its llr_{vg} or its p_{vg} as the observed statistic, because the likelihood ratio and its p -value are two equivalent entities for two reasons. First, the likelihood ratio of each variant v has the same distribution and degree of freedom. Second, the p -value function is 1-to-1 and strictly monotone. Thus, having a null density for the $\max_{v \in V} \{\text{llr}_{vg}\}$ is equivalent to having a null density for the minimum p_g .

We empirically find that on average the correlation of the likelihood ratios at any two variants is roughly equal to their LD; that is, on average $\text{cor}(\text{llr}_{ug}, \text{llr}_{vg}) \approx \text{LD}(u, v)$ for any variant $u, v \in V$ (Fig. 1). For this reason, any function f that accounts for LD and maps an observed test statistic of a gene into an eGene p -value would be applicable in our case. We can use such a function to convert the observed test statistics p_g at the gene g to eGene p -value α_{p_g} without doing the permutation test. Each gene g has its own LD structure and requires its own function f , because the cis-variants of each gene are non-identical.

We apply MVN-EGENE to estimate the function f for each gene. MVN-EGENE is a software that tests if a gene is an eGene in one tissue. MVN-EGENE is designed so that one does not need to do the permutation test when estimating an eGene p -value. MVN-EGENE is unable to simultaneously consider more than one tissues, as would a meta-analysis would.

To compute the function f at a gene, we apply MVN-EGENE at that gene in a tissue (Sul et al., 2015). We assume that the LD does not change much between tissues, and it does not matter much which particular tissue is chosen, as long as it has many samples.

In MVN-EGENE, the test statistic for a gene is the most significant effect size taken over all cis-variants. The p -value of this test statistic depends on the LD of the cis-variants. Instead of doing a permutation test to compute this p -value, MVN-EGENE simulates data under the null hypothesis using a multivariate normal

distribution. In brief, in one simulation, MVN-EGENE samples the effect sizes of the cis-variants of a gene in a tissue using zero as the mean effect and LD as the COV matrix. In this simulation, the most significant effect among these effect sizes is taken to be the test statistic at the gene. After many simulations, one can create a null distribution for the observed test statistic. One can easily convert an effect size into a p -value using a normal distribution. By having a null density of the most significant effect size taken over all the variants, one also has the null density of the minimum p -value taken over all the variants. This null density of the minimum p -value in MVN-EGENE properly handles LD at the gene. Here, we use this distribution of minimum P -values as our null density to convert the observed minimum likelihood ratio p -value p_g to its eGene P -value α_{p_g} in both RECOV and RE2.

2.4.2 Estimating genomic control

In the GTEx dataset and other multi-tissue gene expression datasets, the same individual may provide samples for many tissues (Fig. 2). Sharing of samples from the same individuals among tissues is known to inflate the FPR in a meta-analysis (Han et al., 2016). Before testing whether the RECOV outcome is affected by the fact that tissues share individuals, we test if RECOV inflates the FPR when the data is absolutely free of any spurious statistical association. These signals can be due to LD, shared individuals in tissues, batch effects, or correlated expressions of the same gene (or between genes) across tissues. It is important to mention that in the real GTEx data, batch effects have been dealt with by the GTEx consortium by applying PEER factors on the gene expression in each tissue (The GTEx Consortium, 2015). Subsection 2.4.1 above describes how RECOV and RE2 handle LD in the variants. We now describe how we use a genomic control (GC) factor to remove the effect of shared individuals in the tissues from the meta-analysis results. This GC factor is clearly data dependent as different datasets will require different GC values.

Here, we focus on finding the GC factor for the GTEx data. To do this, we simulate two types of datasets and compare their behaviors. The first type does not contain any spurious statistical signals. The second type contains only signal due to sharing of samples among the tissues, and the number of people shared between pairs of tissues is taken from the GTEx data. Our goal is to apply RECOV and RE2 to the GTEx data; to avoid data reusing, the SNPs and the gene expressions in both types of datasets are simulated and thus are independent of the values in the GTEx data.

When there is not any spurious statistical association in the data, any alternative hypothesis must be rejected more often than the null hypothesis. We simulate data to demonstrate that RECOV does not inflate FPR in this case. In each simulated dataset, the number of individuals per tissue is taken from the GTEx data, but we do not let tissues share individuals. We generate 1000 SNPs at various minor allele frequency (MAF) without LD, and a random gene expression in each tissue. We generate gene expressions where the expression of the same gene is not correlated between any two tissues. We compute the p -value of the likelihood ratio at each SNP using both RECOV and RE2 model. We repeat this simulation 1000 times to obtain 1 000 000 p -values each for RECOV and RE2. The histograms of these p -values in both RECOV and RE2 indicate that the null hypothesis is more favored than the alternative hypothesis (Fig. 3A and B).

To measure the effect strictly caused by shared individuals, we simulate datasets as above, but now allow tissues to share individuals. The number of people shared between pairs of tissues is taken




Fig. 1. Correlation for the likelihood ratios of a pair of cis-variants versus their LD. Denote $\text{cor}(\text{llr}_u, \text{llr}_v)$ as the correlation for the likelihood ratios of variants u and v over all genes where both are cis-SNPs. Empirically, $\text{cor}(\text{llr}_u, \text{llr}_v)$ is close to the LD of u and v . To show this, we randomly select many pairs of cis-SNPs from the gene ENGS00000204219.5 that also appear together in at least two other genes. These pairs are then grouped into bins by their LD (bin width 0.05). We compute the likelihood ratio for each SNP in each pair over all the genes in which they are cis-variants. Using these likelihood ratios, we estimate $\text{cor}(\text{llr}_u, \text{llr}_v)$ for the pair u, v . We average $\text{cor}(\text{llr}_u, \text{llr}_v)$ over all pairs u, v in each LD bin. We then plot the absolute value of this average against the LD value. The identity line is shown in red. Plots for additional pairs chosen from other genes are shown in Supplementary Figure S1




Fig. 2. Shared individuals among the 44 tissues in the GTEx dataset. Degree of sample sharing between two tissues is measured using the Jacquard index

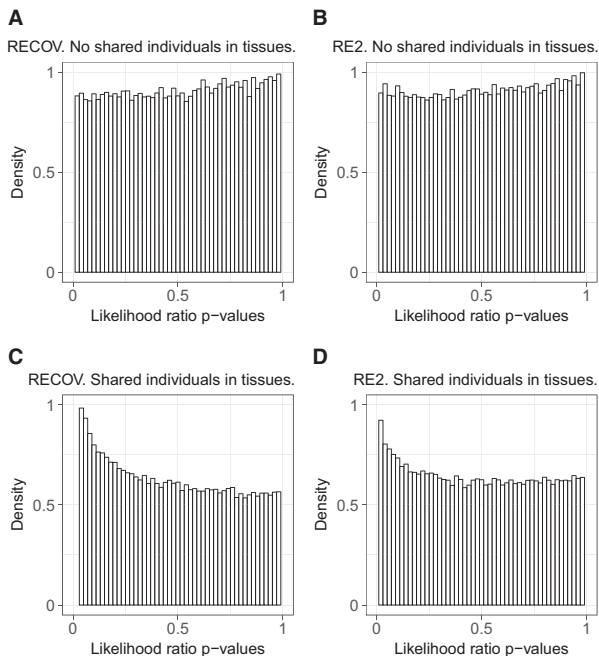


Fig. 3. (A) RECOV and (B) RE2 applied to datasets where the tissues do not share individuals. (C) RECOV and (D) RE2 applied to datasets where the tissues share individuals

from the GTEx data. In each simulation, we compute the likelihood ratio p -values at 1000 SNPs, and repeat the simulation 1000 times to obtain 1 000 000 P -values. We observe that these p -values shift toward 0 when the tissues share samples that are from the same individuals (Fig. 3C and D). In this case, we estimate the RECOV and RE2 GC factor to be 1.2947 and 1.1045, respectively. These GC factors are used to remove the effect caused by shared individuals in tissues that may inflate the FPR. To compute a GC factor, one converts the median of the observed p -values into a chi-square statistic, then finds a multiplying factor to scale this new statistic to a chi-square random variable that has p -value at 0.50 (Devlin and Roeder, 1999).

3 Results

3.1 RECOV controls FPR

When using any meta-analysis method to find an eGene, one needs to apply the method at every cis-variant of the specified gene in order to determine if that gene has at least one eQTL. Thus, the global FPR of RECOV and RE2 depends on the FPR at a single cis-variant. For this reason, we measure the FPR of RECOV and RE2 at a single variant.

To obtain the FPR at one variant, we simulate 1000 datasets for a single variant under the null hypothesis where the variant is not associated with the gene expression in any of 44 tissues. The MAF of the SNP is randomly chosen and kept the same in all 1000 datasets. Then in each dataset, the genotype for this SNP and the gene expression are simulated independently of the values in the GTEx data.

To make the simulated data more realistic, we first let each tissue have the same number of individuals as in the GTEx data, and each pair of tissues have the same number of shared samples as in the GTEx data. Second, we set expression levels of the same gene from the same individual to be correlated with an average correlation of 0.5 across tissues, using the sampling method described in (Sul et al., 2013). This correlation of expression can occur when the tissues of an individual have been exposed to the same environmental factors.

In each of the 1000 datasets, we estimate the effect size and variance of this single variant on the gene expression in each tissue. RECOV and RE2 take these effect sizes and variances and produce a meta-analysis p -value for this variant. The GC factor estimated in Subsection 2.4.2 is used to transform this p -value in each simulation. This removes only the effect of shared individuals, which is not explicitly modeled in RECOV and RE2. The FPR of this single variant is the fraction of times its transformed p -values are significant.

We repeat this experiment for 1000 independent variants, so that we have 1000 measures of FPR for RECOV and RE2. We use the significance level of 0.05 ($\alpha = 0.05$). We find that RECOV attains correct FPR for the majority of variants tested. In RECOV, the median FPR among the 1000 variants is 0.05, and the 75 and 95% quantiles are 0.06 and 0.09. In RE2, the median FPR is 0.05, and the 75 and 95% quantiles are 0.07 and 0.10. These results demonstrate that RECOV and RE2 control the FPRs in a realistic setting.

3.2 RECOV discovers more eGenes in GTEx data

We apply RECOV, RE2 and TBT to the real multi-tissue eQTL dataset from the GTEx consortium. We use GTEx Pilot Dataset V6 released on 12 January 2015. The GTEx consortium has performed RNA-seq on 44 tissues from hundreds of individuals, and we select 15 336 genes that have expression data in all 44 tissues. The consortium has already applied PEER factors to every gene expression in each tissue to remove any batch effects (The GTEx Consortium, 2015). For genotype data, we use the GTEx imputed genotype data that contains 5 million SNPs for each individual. Like in the original GTEx pilot study, for each gene, we use its cis-SNPs, which are defined to be located within 1Mb from its transcription start site (The GTEx Consortium, 2015). Not all variants are genotyped in every tissue, because the 44 tissues contain samples from different individuals. We use only cis-variants that are genotyped in all 44 tissues. The median number of cis-variants tested per gene is 3744.

For each of the 15 336 genes, we apply RECOV, RE2, and TBT to every cis-SNP. For each cis-SNP of a gene, our test statistic is the log likelihood ratio (for RECOV and RE2) or SNP-effect (for TBT). These test statistics are converted into p -values by using a chi-square distribution (for RECOV and RE2) or normal distribution (for

TBT). These p -values are then transformed using the GC factors to remove the effect of shared individuals in the tissues. Finally, the most significant p -value among all cis-variants is converted into an eGene p -value by method in 2.4.1 (for RECOV and RE2) or by EGENE-MVN (for TBT).

After computing the eGene p -values for 15 336 genes, we use Bonferroni correction to control for multiple testing correction at 5% level to identify significant eGene p -values; thus, each gene has a significance threshold of 0.05/15 336.

Figure 4A shows the Venn diagram of the numbers of eGenes found by TBT, RE2 and RECOV. The majority of tested genes are found to be candidate eGenes. This is expected because there are many tissues tested. It is likely that a gene contains at least one eQTLs in some tissue, which significantly increases the total number of eGenes detected. Both RE2 and RECOV find more candidate eGenes than TBT. This result agrees with previous findings where applying meta-analysis to multi-tissue datasets yields better outcome than the simple TBT approach (Flutre *et al.*, 2013; Sul *et al.*, 2013).

RECOV detects the highest number of eGenes among the three methods. Out of the 15 336 genes tested, RECOV finds that 81.40% of those genes are eGenes while TBT and RE2 find 61.90 and 78.45% of genes are eGenes, respectively. This shows that our approach detects 3% more eGenes than RE2 and about 20% more eGenes than TBT.

Next, we apply each method to a case study in order to understand the circumstances where one method outperforms the other two. We begin with the simple TBT method. In Figure 4A, there are 252 genes detected only by the TBT method. Previous publications have reported TBT to be the most powerful option to detect genes with eQTLs that are found in only one tissue (Sul *et al.*, 2013; The GTEx Consortium, 2015). In the TBT method, one analyzes each tissue independently, and is able to determine the number of tissues in which a gene is an eGene. In our result, out of these 252 genes, 225 are eGenes in only 1 tissue, 25 are in 2 tissues, and only 2 are in 3 tissues. This finding agrees with Figure 2 in Sul *et al.* (2013).

Of the 452 genes discovered by only RECOV, the average RECOV eGene p -value is $8.52E^{-9} (\pm 1.51E^{-8})$; whereas the average RE2 eGene p -value is $4.18E^{-3} (\pm 2.85E^{-2})$. To understand why RECOV discovers genes that are not found by TBT and RE2, consider the protein-coding gene CABLES1 (Ensembl id ENSG00000134508.8) which is only detected by RECOV. From the GTEx portal, CABLES1 is expressed mostly in brain tissues, yet it does not have any brain-specific eQTLs. RECOV is a meta-analysis method that pools samples across tissues to increase signals of eQTLs. Thus, when the sample size per tissue is small enough that

eQTL signals may be undetected, RECOV outperforms TBT. Unlike RE2, the meta-analysis of RECOV considers correlation of the cis-variants across the tissues; thus RECOV would be better than RE2 if CABLES1 has a consistent correlation pattern. This is indeed the case (Fig. 4B). CABLES1's RECOV and RE2 eGene p -value are $4.94E^{-13}$ and $5E^{-5}$, respectively.

Of the 88 genes discovered by only RE2, the average RE2 eGene p -value is $1.15E^{-8} (\pm 1.44E^{-8})$; whereas the average RECOV eGene p -value is $1.85E^{-4} (\pm 2.32E^{-4})$. We suspect that these 88 genes are genes with eQTLs in multiple tissues. However, due to low sample size, these eQTLs signals may be undetected or do not produce an eGene q -value less than the significance threshold in TBT analysis. As a case study, consider the protein-coding gene GALNT11 (Ensembl id ENSG00000178234.8) which is detected by only RE2. Like CABLES1, GALNT11 is expressed mostly in the brain tissues (The GTEx Consortium, 2015). Unlike CABLES1, GALNT11 has eQTL signals in the frontal cortex brain tissue, but these signals produce an eGene q -value of 0.0189 which is higher than the TBT significance threshold. In this case, a meta-analysis approach is more suitable because it combines data from many tissues to improve the eGene p -value. GALNT11's cis-variants have correlated effect sizes across the brain tissues, but this pattern does not stand out from the rest of the tissues when compared with that of CABLES1 (Fig. 4C). For this reason, GALNT11's RECOV p -value is higher than its RE2 p -value ($3.50E^{-4}$ versus $7.08E^{-8}$). RE2 may also have better performance than RECOV in cases where the cis-variants do not have an obvious correlation pattern across the 44 tissues. As an example, consider the pseudogene RP11-34P13.16 (Ensembl id ENSG00000269981.1), which is not tissue-specific (The GTEx Consortium, 2015). The effect sizes of its cis-variants appear to be randomly correlated (Fig. 4D), and its RECOV and RE2 p -values are $1.50E^{-4}$ and $1.37E^{-8}$, respectively. Altogether, these attributes may have caused the different results produced by RECOV and RE2.

4 Discussion

In this article, we introduce a new REs meta-analysis method named RECOV. Our approach is motivated by the insight that the same SNP may have similar effect on the same gene in related tissues. We explicitly model these phenomena by adding a COV matrix to the existing RE2 model introduced by Han and Eskin (2011). When applied to the GTEx data, RECOV controls the FPR at the SNP level. More importantly, using no additional data, RECOV finds 3% more eGenes than the TBT and RE2 methods.

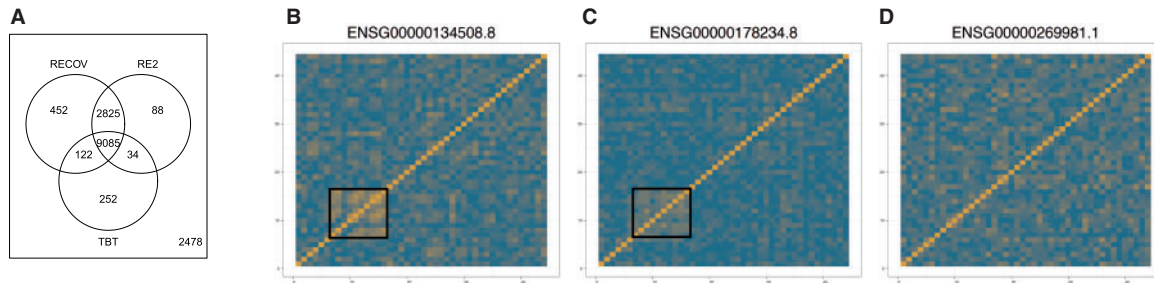


Fig. 4. (A) Venn diagram of the numbers of eGenes found by TBT, RE2 and RECOV. (B) The correlation of SNP-effects for the gene ENSG00000134508.8 in 44 tissues (tissue names are omitted). The correlation is computed by using the matrix B_g in Subsection 2.3.2 where the formula is $B_g B_g^T$ (after proper scaling and removal of nearby SNPs). Black box indicates the brain tissues. ENSG00000134508.8 is found to be an eGene by only the RECOV method. The correlation of SNP-effects for gene (C) ENSG00000178234.8 and (D) ENSG00000269981.1 in 44 tissues (tissue names are omitted). ENSG00000178234.8 and ENSG00000269981.1 are found to be eGenes by only the RE2 method

RECOV scales well to large numbers of tissues compared with previous meta-analysis methods for gene expression data. For example, Meta-Tissue and eQTLBma can only handle up to 10 and 20 tissues, respectively (Flutre *et al.*, 2013; Sul *et al.*, 2013). RECOV also requires only the summary statistics for the SNP effect on the gene expression in each tissue. These summary statistics are often readily available in gene expression data. Thus, unlike the model by Acharya *et al.* (2016), RECOV requires minimal data preparation.

RECOV, and the RE2 it extends, require optimizing two parameters in the log likelihood ratio. These unknowns are the mean effect size and the scaling factor for the COV matrix, both of which can be estimated using efficient heuristic methods. We note that the TBT method avoids this optimization. This is a speed-performance trade-off. This study and others show that the meta-analysis approach is better than TBT when applied to multi-tissue data (Acharya *et al.*, 2016; Flutre *et al.*, 2013; Sul *et al.*, 2013). Unlike TBT, RECOV does not provide information about the specific subset of tissues in which the gene is an eGene. This problem is inherent to all meta-analysis methods, which only test whether a gene is an eGene in at least one tissue.

Next, we address our use of the GC factor for RECOV. The GC factor is traditionally used to correct for inflation due to population structure in classic GWAS, but in this paper, we use it to correct for inflation from any unmodeled source. We show that this inflation is due to tissues containing samples from the same individuals. This problem of sample sharing is not the same as the problem of population structure in GWAS (Han and Eskin, 2011, 2012). The value of the GC factor depends on the choice of the COV matrix U_{vg} in the model. As shown in this article, when $U_{vg} = I$ for RE2 the GC factor is 1.1045, whereas when $U_{vg} = B_{-vg} B_{-vg}^T$ for RECOV the GC factor is 1.2947.

RECOV is a general framework for meta-analysis that can be used with any COV matrix. The COV matrix used in this article (described in Subsection 2.3.2) reflects our assumptions about the behavior of the same SNP in different tissues. Namely, we assume a SNP has correlated effects on a gene's expression across tissues. There are many ways to select this COV matrix, and other options may better fit different assumptions about the data. For example, if we instead assume the same SNP has correlated effects on the expressions of different genes across the tissues, we can estimate U_{vg} by combining information from neighboring genes of g , using knowledge from a gene-gene interaction network. The problem of selecting the most suitable COV matrix for RECOV is a rich topic for future work.

Funding

D.D., L.G., E.K. and E.E. are supported by National Science Foundation [grants 0513612, 0731455, 0729049, 0916676, 1065276, 1302448, 1320589 and 1331176], National Institutes of Health [grants K25-HL080079, U01-DA024417, P01-HL30568, P01-HL28481, R01-GM083198, R01-ES021801, R01-MH101782 and R01-ES022282] and NINDS Informatics Center for Neurogenetics and Neurogenomics [grant P30 NS062691]. B.H. is supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIP) [grant 2016R1C1B2013126]. E.E. is supported in part by the NIH BD2K award [grant U54EB020403]. D.D. is supported by the NIH Training Grant in Genomic Analysis and Interpretation [grant T32HG002536].

Conflict of Interest: none declared.

References

- Abraham, B. and Ledolter, J. (2006). *Introduction to regression modeling*. Thomson Brooks/Cole, Belmont, CA, Thomson Brooks/Cole.
- Acharya, C.R. *et al.* (2016) Exploiting expression patterns across multiple tissues to map expression quantitative trait loci. *BMC Bioinformatics*, **17**, 257.
- Albert, F.W. (2016) Brains, genes and power. *Nat. Neurosci.*, **19**, 1428–1430.
- Dabney, A. *et al.* (2010). qvalue: Q-value estimation for false discovery rate control. *R package version*, 1(0).
- Darnell, G. *et al.* (2012) Incorporating prior information into association studies. *Bioinformatics*, **28**, i147–i153.
- Devlin, B. and Roeder, K. (1999) Genomic control for association studies. *Biometrics*, **55**, 997–1004.
- Duong, D. *et al.* (2016) Using genomic annotations increases statistical power to detect eGenes. *Bioinformatics*, **32**, i156–i163.
- Eskin, E. (2015) Discovering genes involved in disease and the mystery of missing heritability. *Commun. ACM*, **58**, 80–87.
- Flutre, T. *et al.* (2013) A statistical framework for joint eQTL analysis in multiple tissues. *PLoS Genet.*, **9**, e1003486.
- Fromer, M. *et al.* (2016) Gene expression elucidates functional impact of polygenic risk for schizophrenia. *Nat. Neurosci.*, **19**, 1442–1453.
- Han, B. and Eskin, E. (2011) Random-effects model aimed at discovering associations in meta-analysis of genome-wide association studies. *Am. J. Hum. Genet.*, **88**, 586–598.
- Han, B. and Eskin, E. (2012) Interpreting meta-analyses of genome-wide association studies. *PLoS Genet.*, **8**, e1002555.
- Han, B. *et al.* (2009) Rapid and accurate multiple testing correction and power estimation for millions of correlated markers. *PLoS Genet.*, **5**, e1000456.
- Han, B. *et al.* (2016) A general framework for meta-analyzing dependent studies with overlapping subjects in association mapping. *Hum. Mol. Genet.*, **25**, 1857–1866.
- Hormozdiari, F. *et al.* (2015) Identification of causal genes for complex traits. *Bioinformatics*, **31**, i206–i213.
- Huang, Y.-T. *et al.* (2014) Joint analysis of snp and gene expression data in genetic association studies of complex diseases. *Ann. Appl. Stat.*, **8**, 352.
- Joo, J.J. *et al.* (2014) Effectively identifying regulatory hotspots while capturing expression heterogeneity in gene expression studies. *Genome Biol.*, **15**, r61.
- Joo, J.W.J. *et al.* (2016) Multiple testing correction in linear mixed models. *Genome Biol.*, **17**.
- Kang, E.Y. *et al.* (2014) Meta-analysis identifies gene-by-environment interactions as demonstrated in a study of 4,965 mice. *PLoS Genet.*, **10**, e1004022.
- Liu, G. *et al.* (2016) Cis-eQTLs regulate reduced LST1 gene and NCR3 gene expression and contribute to increased autoimmune disease risk: Table 1. *Proc. Natl. Acad. Sci. USA*, **113**, E6321–E6322.
- Nieuwenhuis, M.A. *et al.* (2016) Combining genomewide association study and lung eQTL analysis provides evidence for novel genes associated with asthma. *Allergy*, **71**, 1712–1720.
- Self, S.G. and Liang, K.Y. (1987) Asymptotic properties of maximum likelihood estimators and likelihood ratio tests under nonstandard conditions. *J. Am. Stat. Assoc.*, **82**, 605–610.
- Sul, J.H. *et al.* (2013) Effectively identifying eQTLs from multiple tissues by combining mixed model and meta-analytic approaches. *PLoS Genet.*, **9**, e1003491.
- Sul, J.H. *et al.* (2015) Accurate and fast multiple-testing correction in eQTL studies. *Am. J. Hum. Genet.*, **96**, 857–868.
- The GTEx Consortium. (2015) The genotype-tissue expression (GTEx) pilot analysis: Multitissue gene regulation in humans. *Science*, **348**, 648–660.
- Thompson, S.G. and Sharp, S.J. (1997) Explaining heterogeneity in meta-analysis: A comparison of methods. *Stat. Med.*, **18**, S82.